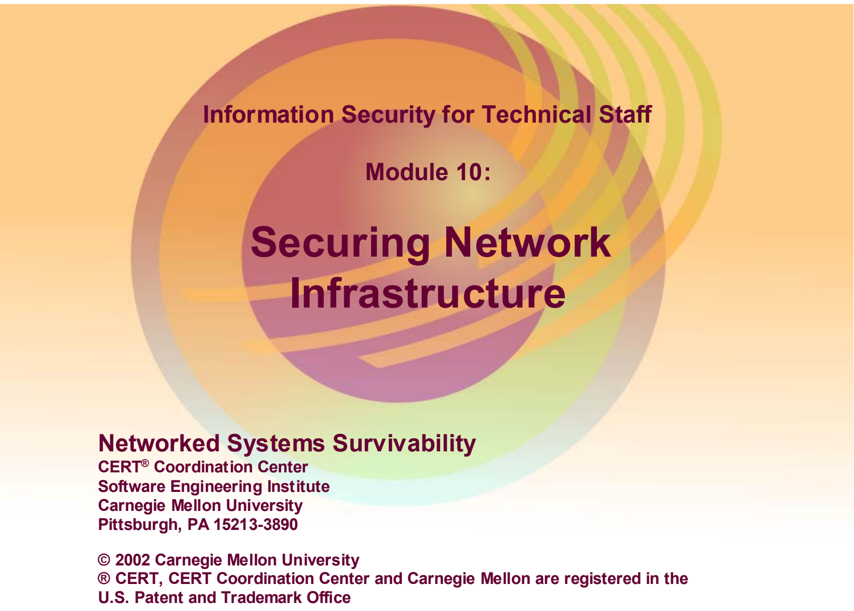

Information Security for Technical Staff

Module 10:

Securing Network Infrastructure

Networked Systems Survivability
CERT® Coordination Center
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213-3890

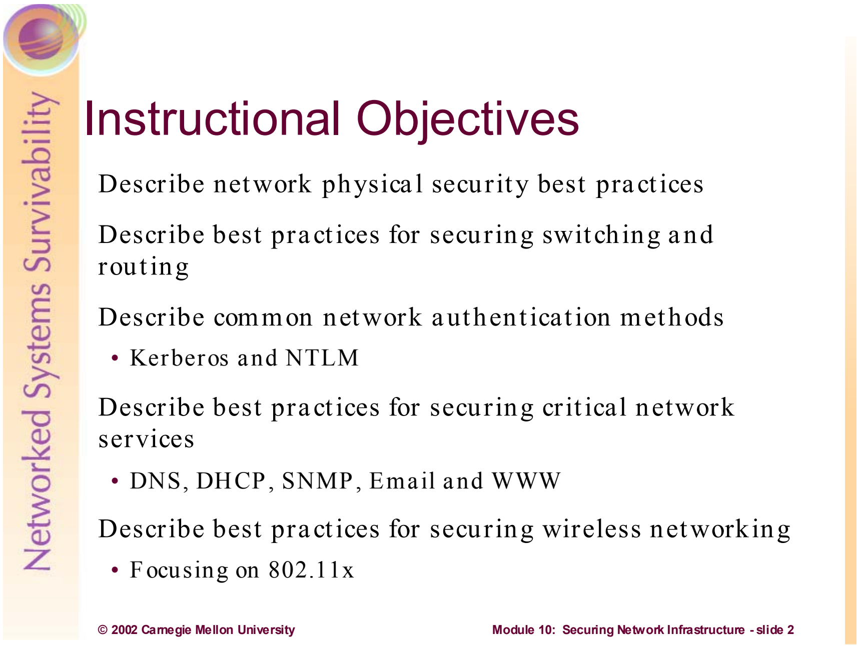

Networked Systems Survivability

Instructional Objectives

Describe network physical security best practices

Describe best practices for securing switching and routing

Describe common network authentication methods

- Kerberos and NTLM

Describe best practices for securing critical network services

- DNS, DHCP, SNMP, Email and WWW

Describe best practices for securing wireless networking

- Focusing on 802.11x

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 2

This module makes the assumption that students have some background knowledge of switching, routing, and key network services. As a result, it doesn't cover all of the basics of what these systems/services do in a network. Rather, it focuses on how best to secure them.

While some attempt is made to keep this module "vendor unspecific," examples will be used to reinforce concepts. Specifics on Cisco IOS based products are used extensively for the routing and switching portions of this module. Similarly, Microsoft Windows NT and 2000 are used for examples in other portions of the module.

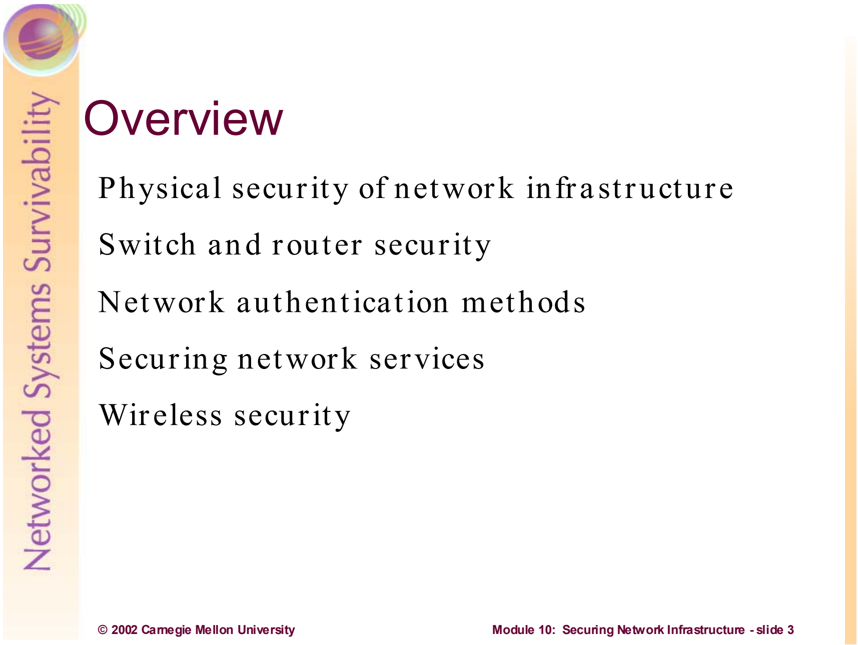

Networked Systems Survivability

Overview

Physical security of network infrastructure

Switch and router security

Network authentication methods

Securing network services

Wireless security

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 3

These topics will be covered in this module.

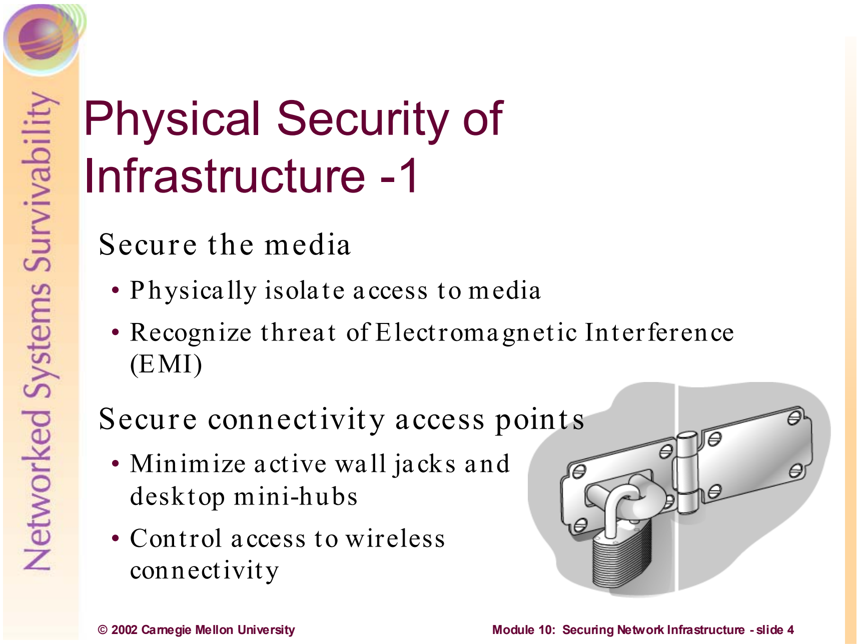

Networked Systems Survivability

Physical Security of Infrastructure -1

Secure the media

- Physically isolate access to media
- Recognize threat of Electromagnetic Interference (EMI)

Secure connectivity access points

- Minimize active wall jacks and desktop mini-hubs
- Control access to wireless connectivity

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 4

Accessibility of transmission media (wiring) is frequently overlooked during security reviews. It is not difficult to intercept data packets by passively tapping exposed wiring. Equipment to do just this is inexpensive and has been used by power and telecommunications companies for years. Electromagnetic interference is electrical energy that emanates from unshielded media. Some media types are more vulnerable to EMI than others. The most popular media used today, category 5 unshielded twisted pair is particularly susceptible to EMI leakage. As the name states, the media is unshielded, with wires enclosed only in plastic sleeves. An extra layer of shielding called plenum is available for category 5 cable, however it is more expensive and not widely implemented. Fiber Optic cabling is immune to EMI altogether and is widely implemented as a backbone media. Keeping network media physically isolated is an excellent practice from a security standpoint.

It is also critical to limit active network entry points. Unused wall jacks for network connectively should be disabled. Widespread use of desktop mini-hubs should be limited, as they provide unmonitored access to the network. Keep charts in wiring closets for which hub/switch ports are activated and monitor for variance.

Wireless connectivity should be managed carefully and procedures should be established for controlling access. This will be discussed later in this module.

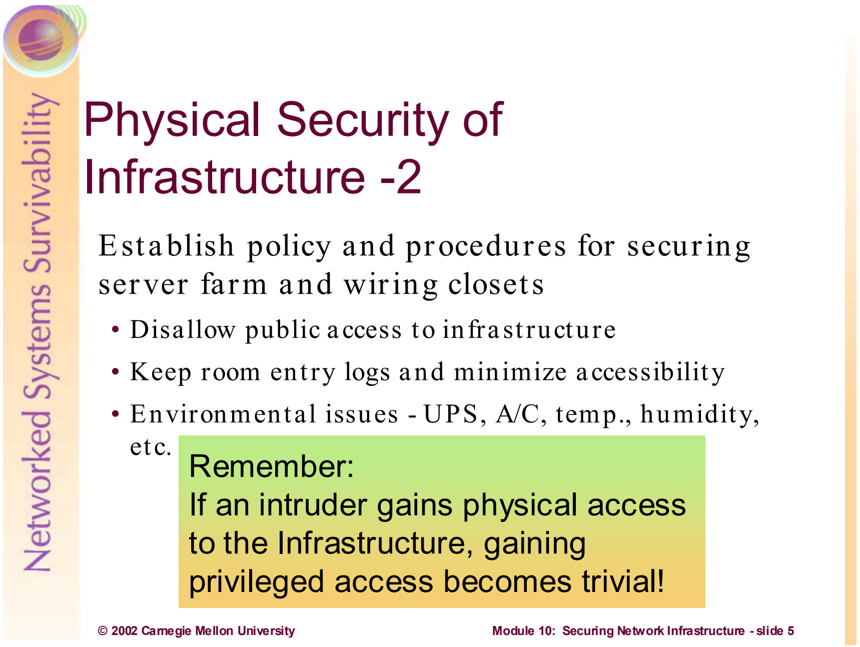
Networked Systems Survivability

Physical Security of Infrastructure -2

Establish policy and procedures for securing server farm and wiring closets

- Disallow public access to infrastructure
- Keep room entry logs and minimize accessibility
- Environmental issues - UPS, A/C, temp., humidity, etc.

Remember:
If an intruder gains physical access to the Infrastructure, gaining privileged access becomes trivial!

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 5

Once an individual has physical access to a piece of networking equipment there is no way to stop him/her from modifying the system. This problem is not only confined to network devices but is also true of computers and any other electrical or mechanical device. It is always a matter of time and effort. There are things that can be done to make this more difficult, but a knowledgeable attacker with access can never be completely defeated, only slowed down. One of the best additions to the security features of a computer network is to limit access. Network infrastructure components, like routers, are especially important because they are often used to protect segments of the network and can also be used for launching attacks against other network segments [NSA1].

Because of the critical nature of network infrastructure, it is essential to physically safeguard it and take measures to increase its overall survivability. Following are best practices for doing so:

- **Maximize structural protection:** A secure room should have full height walls and fireproof ceilings.
- **Minimize external access (doors):** A secure room should only have one or two doors--they should be solid, fireproof, lockable, and observable by assigned security staff. Doors to the secure room should never be propped open.
- **Minimize external access (windows):** A secure room should not have excessively large windows. All windows should have locks.
- **Maintain locking devices responsibly:** Locking doors and windows can be an effective security strategy as long as appropriate authorities maintain the keys and combinations responsibly. If there is a breach, each compromised lock should be changed.
- **Investigate options other than traditional keyhole locks for securing areas as is reasonable:** Based on the findings from your risk assessment consider alternative physical security strategies such as window bars, anti-theft cabling (i.e., an alarm sounds when any piece of equipment is disconnected from the system), magnetic key cards, and motion detectors. Have physical locks as backups for key card systems.
- **Keep a record of your equipment:** Maintain up-to-date logs of equipment manufacturers, models, and serial numbers in a secure location. Be sure to include a list of all attached peripheral equipment. Consider videotaping the equipment (including close-up shots) as well. Such clear evidence of ownership can be helpful when dealing with insurance companies.

- **Make unauthorized tampering with equipment difficult:** Replace regular body case screws with Allen-type screws or comparable devices that require a special tool (e.g., an Allen wrench) to open them.
- **Limit and monitor access to equipment areas:** Keep an up-to-date list of personnel authorized to access sensitive areas. Never allow equipment to be moved or serviced unless the task is pre-authorized and the service personnel can produce an authentic work order and verify who they are. Require picture or other forms of identification if necessary. Logs of all such activity should be maintained. Staff should be trained to always err on the cautious side (and the organization must support such caution even when it proves to be inconvenient).
- **Maintain a reasonable climate within the room:** A good rule of thumb is that if people are comfortable, then equipment is usually comfortable--but even if people have gone home for the night, room temperature and humidity cannot be allowed to reach extremes (i.e., it should be kept between 50 and 80 degrees Fahrenheit and 20 and 80 percent humidity). Note that it's not freezing temperatures that damage disks, but the condensation that forms when they thaw out.
- **Maintain and repair equipment:** Have plans in place for emergency repair of critical equipment. Either have a technician who is trained to do repairs on staff or make arrangements with someone who has ready access to the site when repair work is needed. If funds allow, consider setting up maintenance contracts for your critical equipment. Local computer suppliers often offer service contracts for equipment they sell, and many workstation and mainframe vendors also provide such services. Once you've set up the contract, be sure that contact information is kept readily available. Technical support telephone numbers, maintenance contract numbers, customer identification numbers, equipment serial numbers, and mail-in information should be posted or kept in a log book near the system for easy reference. Remember that computer repair technicians may be in a position to access your confidential information, so make sure that they know and follow your policies regarding outside employees and contractors who access your system [NCES].

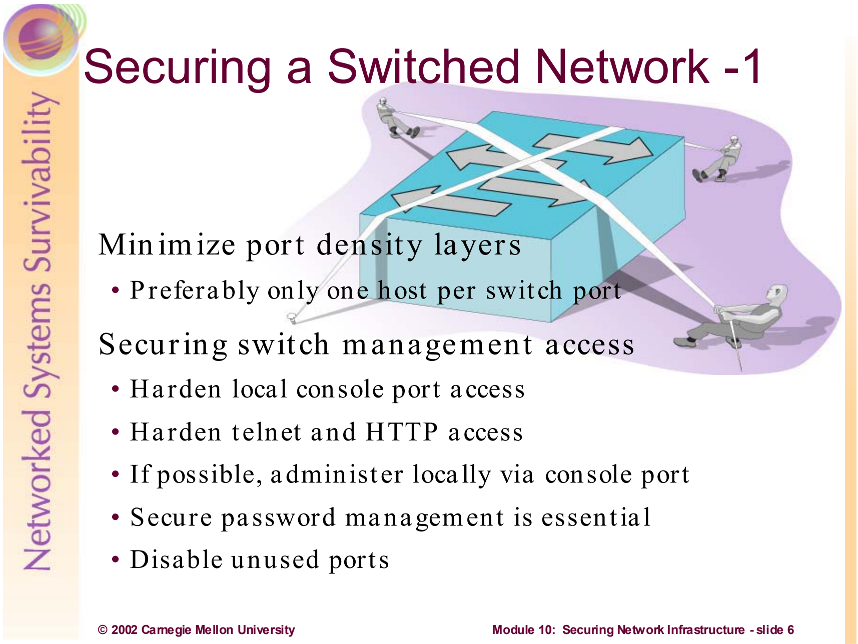

Networked Systems Survivability

Securing a Switched Network -1

Minimize port density layers

- Preferably only one host per switch port

Securing switch management access

- Harden local console port access
- Harden telnet and HTTP access
- If possible, administer locally via console port
- Secure password management is essential
- Disable unused ports

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 6

Switches are a fundamental part of most networks. They make it possible for several users to send information over a network at the same time without slowing each other down. Just like routers allow different networks to communicate with each other, switches allow different nodes (a network connection point, typically a computer) of a network to communicate directly with one another in a smooth and efficient manner.

An illustration of a Cisco Catalyst switch.

There are a lot of different types of switches and networks. Switches that provide a separate connection for each node in a company's internal network are called LAN switches. Essentially, a LAN switch creates a series of instant networks that (should) contain only the two devices communicating with each other at that particular moment [Tyson].

Port Density is defined as the number of ports on a device, such as a network switch, router, or hub. The more ports on a switch, the greater its port density. It determines the number of devices that can be supported by the unit. It is a good security and performance practice to not overload the port density as specified by the switch manufacturer. That is to say, don't plug a bunch of hubs into the ports of your switch to give you more PC connections. You would be expanding collision domains by adding layers of hubs—again, should be discouraged. Doing this will increase the potential for collisions, but more importantly, it will make it easier for intruders to sniff data packets on the network. By limiting the total number of hosts per switch port (1 is best), intruders will have to work much harder to gain access to network resources. This practice will also keep switch performance optimal and protect against processing and memory saturation problems.

It is important to secure administrative access to networking equipment. Written policies and procedures should be in place to define this process. Some security best practices are discussed below:

- Local administrative access (plugging a serial cable from your laptop into the console port of the switch) is more secure than via remote access (Telnet, HTTP, or SNMP) and is therefore preferred
- Set passwords and terminal session timeouts on the console port and the virtual terminal (vty) ports--see specific Cisco IOS commands later under Router Hardening
- Disable HTTP access—when enabled, makes it easier for intruders to gain access to the switch. Not all switches come with this feature but regardless, it should be disabled for stronger security--see specific Cisco IOS commands later under Router Hardening
- Use strong passwords and change them regularly
- Disable all unused switch ports. If physical security to the wiring closet is breached, connecting to the network will still be more difficult for an intruder—think defense in depth!

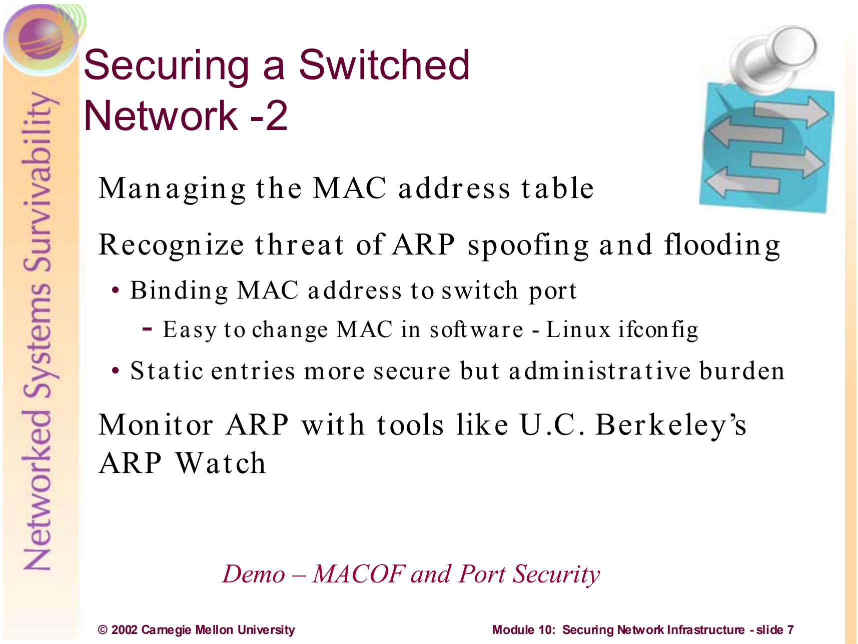

Networked Systems Survivability

Securing a Switched Network -2

Managing the MAC address table

Recognize threat of ARP spoofing and flooding

- Binding MAC address to switch port
  - Easy to change MAC in software - Linux ifconfig
- Static entries more secure but administrative burden

Monitor ARP with tools like U.C. Berkeley's ARP Watch

Demo – MACOF and Port Security

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 7

One of the most challenging tasks when securing a switch is dealing with inherent insecurities associated with the Address Resolution Protocol (ARP)—see Module 5 TCP/IP Security.

The MAC address table on a switch is used to map specific switch ports to MAC addresses. Switches typically build this table dynamically by recording the source MAC address from frames as they pass through switch ports. These entries have timers associated with them, and will be deleted from the table once the timer has expired. Every time a frame arrives at a switch port, the table is consulted to determine if it has an entry that matches the destination MAC address of the frame. If an entry exists, the frame is forwarded only to the mapped port as recorded in the MAC table. If no entry exists, the frame is (typically) broadcasted out to all switch ports. So as you can see, effective management of the MAC table is critical to the operations of a switch.

Many intruder tools exist for tampering with MAC tables and the ARP process[1]. These tool capabilities range from eavesdropping and man-in-the-middle attacks, to Denials of Service. Also, software implementation vulnerabilities abound that make it easy for intruders to wreak havoc on networks.[2]

Some techniques for securing MAC tables are:

- Some switches allow administrators to bind specific MAC addresses to switch ports. The process of dynamically building the MAC table remains the same, however once MAC addresses have been bound to particular switch ports, no other MAC address can be mapped to that port in the MAC table (without the administrator).
- Administrators can also choose to manually build MAC tables for switches—thereby eliminating the dynamic entries, timers, etc. This can be a huge administrative burden and is really only feasible in very small networks.

These techniques are only moderately effective because it is possible to change the MAC address in software or to reconfigure the EEPROM on network cards. Linux users can do this with the ifconfig command and numerous NIC drivers allow you to "soft set" the MAC address. Additionally, it's the TCP/IP stack that actually builds the frames, and there are plenty of tools out there that allow an intruder to craft bogus packets with whatever addresses they want.

[1] http://packetstorm.decepticons.org/papers/protocols/intro_to_arp_spoofing.pdf

[2] http://www.cisco.com/warp/public/707/IOS-arp-overwrite-vuln-pub.shtml

Realizing this network vulnerability should increase the vigilance of administrators towards maintaining strong physical security and locking down access to network connectivity.

It is also recommended that administrators utilize ARP monitoring tools like ARPWatch and ARPSnmp—produced by the University of California's Lawrence Berkeley National Laboratory.[3] These utilities monitor Ethernet traffic and automatically build databases for tracking IP to MAC pairings on your network. They have the capability of emailing administrators when variances in the pairings are detected. This can help to fight MAC spoofing intrusions.[4] These tools are typically included in the standard TCPDump package provided with most distributions of Linux.

[3] http://www.openbsd.org/2.7_packages/i386/arpwatch-2.1a4.tgz-long.html

[4] http://www.csnc.ch/downloads/docs/installation/arpwatch_CSNC.pdf

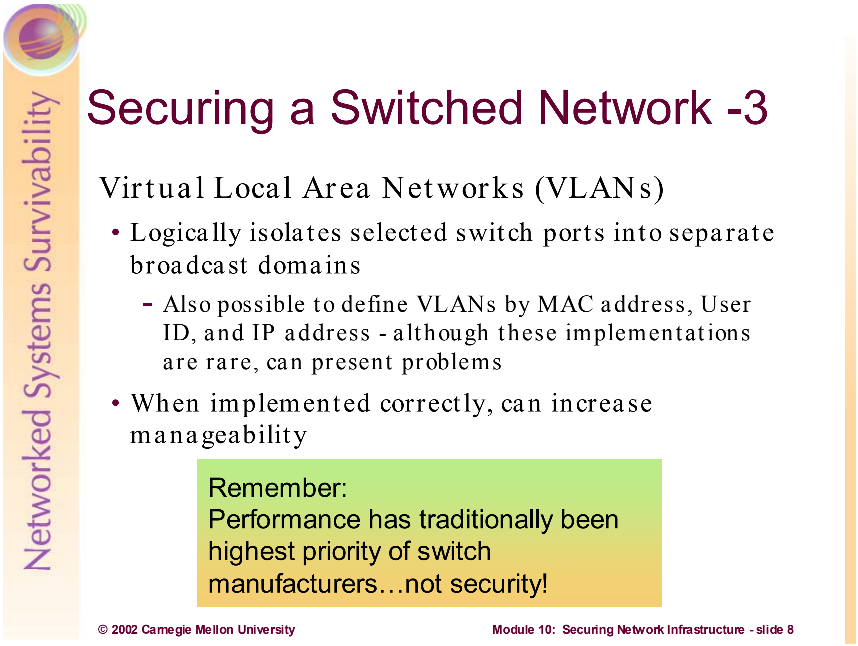

Networked Systems Survivability

Securing a Switched Network -3

Virtual Local Area Networks (VLANs)

- Logically isolates selected switch ports into separate broadcast domains
  - Also possible to define VLANs by MAC address, User ID, and IP address - although these implementations are rare, can present problems
- When implemented correctly, can increase manageability

Remember:
Performance has traditionally been highest priority of switch manufacturers…not security!

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 8

An overview of VLAN technology follows:

A VLAN is a collection of nodes that are grouped together in a single broadcast domain that is based on something other than physical location. A broadcast domain is a network (or portion of a network) that will receive a broadcast packet from any node located within that network. In a typical network, everything on the same side of the router is all part of the same broadcast domain (because routers do not forward broadcasts). A switch that you have implemented VLANs on has multiple broadcast domains, similar to a router. But you still need a router to route from one VLAN to another; the switch can't do this by itself. Here are some common reasons why an organization might have VLANs:

- **Security:** Separating systems with sensitive data from the rest of the network decreases the chance that someone will gain access to information they are not authorized to see. It is important to recognize that VLANs can enhance the security posture of an organization, however they should NOT be viewed as (or implemented as) a stand alone “security solution.”
- **Projects/Special applications**: Managing a project or working with a specialized application can be simplified by the use of a VLAN that brings all of the required nodes together.
- **Performance/Bandwidth**: Careful monitoring of network use allows the network administrator to create VLANs that reduce the number of router hops and increase the apparent bandwidth for network users.
- **Broadcasts/Traffic flow**: Since a principle element of a VLAN is the fact that it does not pass broadcast traffic to nodes that are not part of the VLAN, it automatically reduces broadcasts.
- **Departments/Specific job types:** Organizations may want VLANs set up for departments that are heavy network users (such as Multimedia or Engineering), or a VLAN across departments that is dedicated to specific types of employees (such as managers or sales people).

You can create a VLAN using most switches simply by logging into the switch (via the console port or Telnet) and entering the parameters for the VLAN (name, domain and port assignments). After you have created the VLAN, any network segments connected to the assigned ports will become part of that VLAN.

While you can have more than one VLAN on a switch, they cannot communicate directly with one another on that switch. If they could, it would defeat the purpose of having a VLAN, which is to isolate a part of the network. Communication between VLANs requires the use of a router. VLANs can span across multiple switches and you can have more than one VLAN on each switch. For multiple VLANs on multiple switches to be able to communicate via a single link between the switches, you must use a process called trunking; trunking is the technology that allows information from multiple VLANs to be carried over just one link between switches.

The VLAN Trunking Protocol (VTP) is the protocol that switches use to communicate among themselves about VLAN configuration.

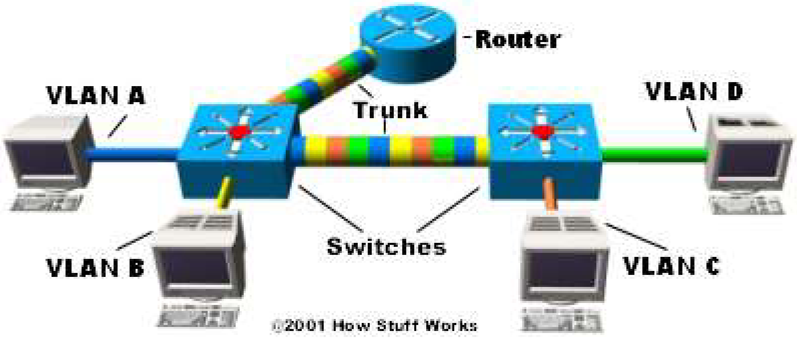

In the image above, each switch has two VLANs. On the first switch, VLAN A and VLAN B are sent through a single port (trunked) to the router and through another port to the second switch. VLAN C and VLAN D are trunked from the second switch to the first switch, and through the first switch to the router. This trunk can carry traffic from all four VLANs. The trunk link from the first switch to the router can also carry all four VLANs. In fact, this one connection to the router allows the router to appear on all four VLANs, as if it had four, different, physical ports connected to the switch.

The VLANs can communicate with each other via the trunking connection between the two switches using the router. For example, data from a computer on VLAN A that needs to get to a computer on VLAN B (or VLAN C or VLAN D) must travel from the switch to the router and back again to the switch. Because of the transparent bridging algorithm[5] and trunking, both PCs and the router think that they are on the same physical segment [Tyson].

VLAN membership can be derived from:

- Switch ports
- MAC addresses of hosts connected to a switch
- The type of Layer 3 protocol (IP, IPX, NetBIOS, etc.)
- User ID

Of these, membership derived by configuring individual switch ports is by far the most common. Utilizing MAC addresses for VLAN membership can be very time consuming to configure and makes it very difficult to have individual hosts be a part of multiple VLANs. Layer 3 membership is flexible because you can segregate different types of networks from each other (i.e. keep IP separate from IPX). However, most IP addresses are distributed via DHCP and are therefore dynamic—this can make for excessive updates and added complexity. Using User ID's for VLAN membership is interesting technology; as it incorporates pieces of all of the above membership schemes. It allows for the greatest flexibility and user mobility within the network, however, it comes with a sizeable overhead of hardware and software and the added complexity of multi-platform interoperability issues.[6]

[5] http://www.howstuffworks.com/lan-switch4.htm
[6] http://www.cisco.com/univercd/cc/td/doc/pcat/wrurto.htm

Another important distinction between VLAN implementations is the method used to indicate membership when a packet travels between switches. Two methods exist — implicit and explicit.

- **Implicit**: VLAN membership is indicated by the MAC address. In this case, all switches that support a particular VLAN must share a table of member MAC addresses.
- **Explicit**: A tag is added to the packet to indicate VLAN membership. Cisco ISL and the IEEE 802.1q VLAN specifications both use this method.

To summarize, when a packet enters its local switch, the determination of its VLAN membership can be port-based, MAC-based or protocol-based. When the packet travels to other switches, the determination of VLAN membership for that packet can be either implicit (using the MAC address) or explicit (using a tag that was added by the first switch). Port-based and protocol-based VLANs use explicit tagging as their preferred indication method. MAC-based VLANs are almost always implicit.

The bottom line is that the IEEE 802.1q specification is going to support port-based membership and explicit tagging, so these will be the default VLAN model in the future.[7]

[7] http://www.intel.com/network/connectivity/resources/doc_library/tech_brief/virtual_lans.pdf

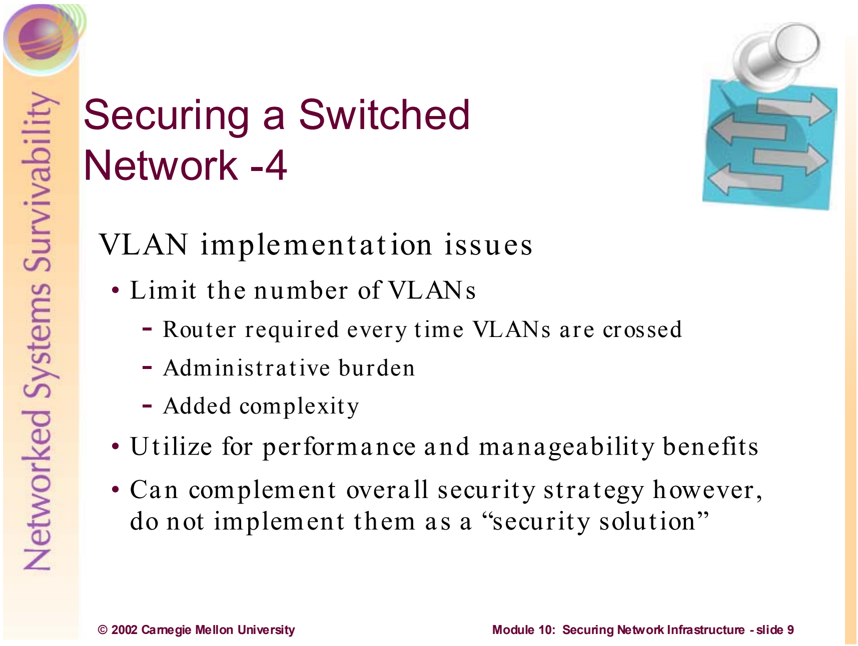

Networked Systems Survivability

Securing a Switched Network -4

VLAN implementation issues

- Limit the number of VLANs
  - Router required every time VLANs are crossed
  - Administrative burden
  - Added complexity
- Utilize for performance and manageability benefits
- Can complement overall security strategy however, do not implement them as a "security solution"

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 9

Using VLANs in your network design can help you solve business and technical needs, but they should be used with discretion. Creating too many VLANs in your network design can cause an administrative nightmare. If your organization is going to invest in a Layer 2 switch that supports VLANs, take advantage of the switching technology. Layer 2 switches provide wire-speed forwarding of frames, and do not incur the latency that traditional software-based routers do. If you are going to build a switched network, try to switch using Layer 2 as much as possible, and route using Layer 3 when necessary.

Remember that each VLAN you create essentially creates a Layer 3 network that must be routed, so unless your traffic is purely workgroup-based, you will always need a routing function in your network. The rapid growth of e-mail, intranets, and the Internet led to the rapid growth of server farms. Server farms can contain shared file, application, and database servers, usually grouped in a dedicated VLAN or VLANs, and require users to communicate across VLAN boundaries using a router. As a reminder, try to keep your design as simple and flexible as possible. Start simple first, then implement a more complex design if requirements can't be met with the existing design. Use VLANs to make your life easier, not more difficult.

The most fundamental benefit of VLAN technology is the capability to create workgroups based on function rather than on physical location or media. Traditionally, network administrators grouped users under the same functional department by physically moving users, their desktop and servers into a common environment such as a shared LAN segment. All team members had to be physically connected to the same media to take advantage of the localized higher-speed server connection. VLANs allow administrators to create, group, and regroup LAN segments logically and instantaneously, without changing physical infrastructure and taking users and servers down. The ability to easily add, move, and change users to the network is a key benefit of VLANs.

VLANs also offer the some measure of security. Users in a defined group are prevented from accessing another group's data, because each VLAN is a closed, logically defined group. Imagine a company in which the Accounting department, working on confidential financial statements, is spread across all three floors of a building. The Engineering and Marketing departments are spread across all three floors as well. Using VLANs, the Engineering and Marketing workgroups can be located on all three floors as members of two different VLANs, and the Accounting department can be members of a third VLAN that spans all three floors. Now the network traffic generated by Accounting will only be accessible to employees of that department, and the Engineering and Marketing teams will not be able to access

Accounting's confidential data. Obviously, there are several other requirements to ensure complete security, but VLANs can be part of an overall network security strategy. The below diagram illustrates how functional VLANs can span traditional physical boundaries.

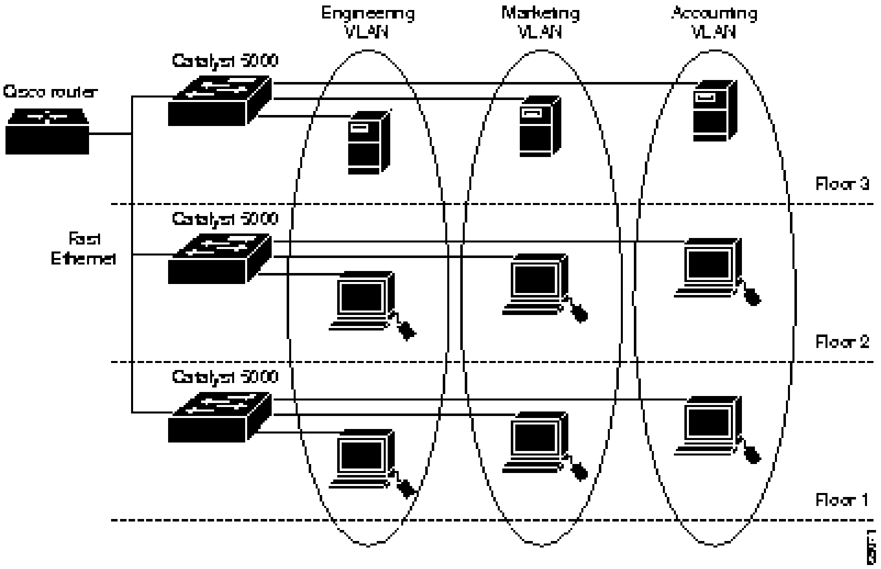

Since VLANs are defined within the device, they can be quickly and easily modified at any time to add, delete, move, or change users as required. This example reinforces the point of utilizing VLANs for the performance and manageability they provide. It also shows how VLANs can enhance security to a degree by logically separating traffic.

As we'll see on the next slide, VLANs have many security concerns associated with them and therefore, (as mentioned previously) they should not be thought of as a "security solution." Rather, they should be implemented as only one of several security measures that contribute to a defense-in-depth strategy.

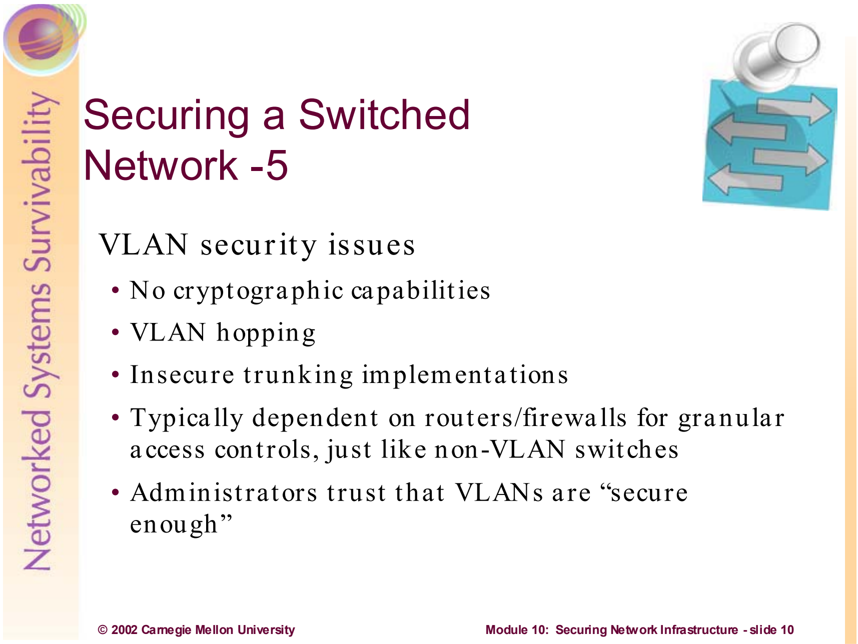

Networked Systems Survivability

Securing a Switched Network -5

VLAN security issues

- No cryptographic capabilities
- VLAN hopping
- Insecure trunking implementations
- Typically dependent on routers/firewalls for granular access controls, just like non-VLAN switches
- Administrators trust that VLANs are "secure enough"

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 10

VLANs are often marketed as "security solutions" by over-zealous (or under-informed) salespeople. Unfortunately, VLAN technology was not designed with security in mind. As a result, it has a number of security issues surrounding it:

- VLAN capable switches typically don't have any cryptographic authentication and encryption capabilities.
- Few implementations of VLAN trunking offer any authentication mechanisms for trunked VLAN communication. That is to say, switches don't check to see if the trunking packets they are receiving (which may be instructing them to change their VLAN configurations) are actually from another switch on the network. It has been proven that the IEEE's 802.1Q VLAN trunking standard as implemented on Cisco switches (and presumably other vendors products as well) is susceptible to attack. Because of a lack of any authentication, forged 802.1q frames can be used to hop between VLANs on the same switch and on other network switches without using a router.[8] Additionally, Cisco's proprietary Inter-Switch Link (ISL) VLAN trunking protocol is set to Auto mode by default.[9] This means it will trust the switch at the other end of the wire to tell it how to handle its VLAN trunking. Again, there is no authentication process built into the specification[10] and is therefore a security risk.
- Another security issue is that most VLAN implementations rely on routers and firewalls to provide any granular access controls. VLANs simply allow or deny all packets based solely on VLAN membership. It doesn't allow you to permit or deny certain protocols or service ports into or out of a VLAN. This seems obvious because switches are layer 2 devices, however it highlights the notion that VLANs are NOT a security technology.
- It's a bad idea to place too much faith in any one technology where security is concerned.

[8] http://www.sans.org/newlook/resources/IDFAQ/vlan.htm

[9] http://lists.synfin.net/Archives/firewall-wizards/1998/Nov/msg00055.html

[10] http://www.cisco.com/warp/public/473/741_4.html

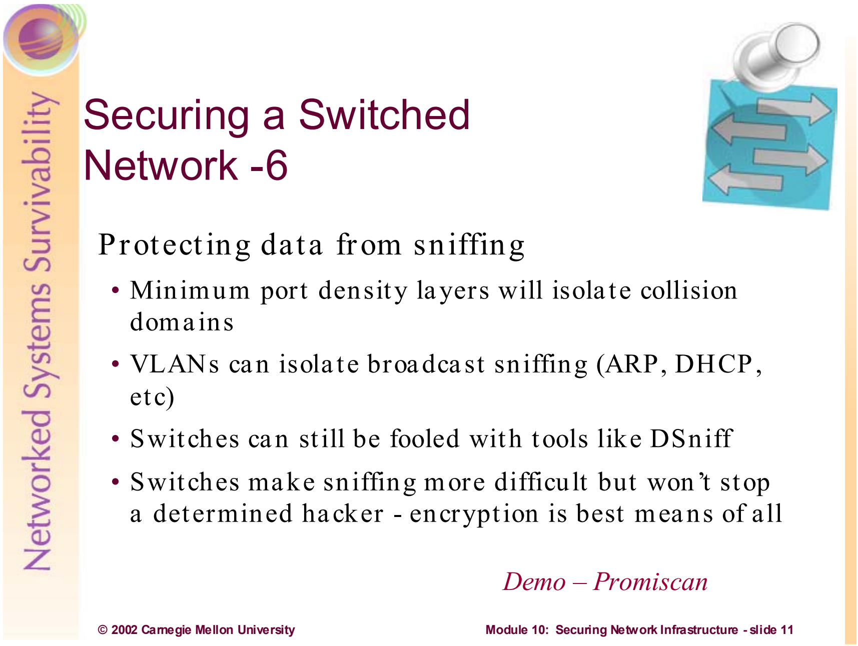

Networked Systems Survivability

Securing a Switched Network -6

Protecting data from sniffing

- Minimum port density layers will isolate collision domains
- VLANs can isolate broadcast sniffing (ARP, DHCP, etc)
- Switches can still be fooled with tools like DSniff
- Switches make sniffing more difficult but won't stop a determined hacker - encryption is best means of all

Demo – Promiscan

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 11

Switching has long been seen (inaccurately) as a cure for network sniffing. By following the best practice of attaching only 1 host per switch port throughout your network, you will keep the size of your collision domains to the minimum. This means that users will only be able to sniff traffic destined for their MAC address (either unicast or broadcast). Implementing VLANs will segment your broadcast domains, and will further limit the amount of broadcast traffic that can be sniffed.

The above strategies should be implemented, however it is not fool proof for stopping unauthorized packet sniffing on networks. There are a number of tools available that take advantage of weak protocols (like ARP) to fool switches into passing all network traffic to an intruder.[11] One in particular, called dsniff, is freely available and is very effective in gaining access to data otherwise blocked by switching. It is available on Linux and Windows platforms and can be used by administrators to monitor their networks—although this should be done with extreme caution.[12]

> Dsniff is a collection of tools for network auditing and penetration testing. dsniff, filesnarf, mailsnarf, msgsnarf, urlsnarf, and webspy passively monitor a network for interesting data (passwords, e-mail, files, etc.). arpspoof, dnsspoof, and macof facilitate the interception of network traffic normally unavailable to an attacker (e.g, due to layer-2 switching). sshmitm and webmitm implement active monkey-in-the-middle attacks against redirected SSH and HTTPS sessions by exploiting weak bindings in ad-hoc PKI.[13]

The bottom line is don't count on switching to solve your data confidentiality concerns. Use encryption technologies (like IPSec or VPNs) for this.

[11] http://rr.sans.org/switchednet/layer2.php

[12] http://www.linuxsecurity.com/feature_stories/feature_story-89-print.html

[13] http://naughty.monkey.org/~dugsong//dsniff/

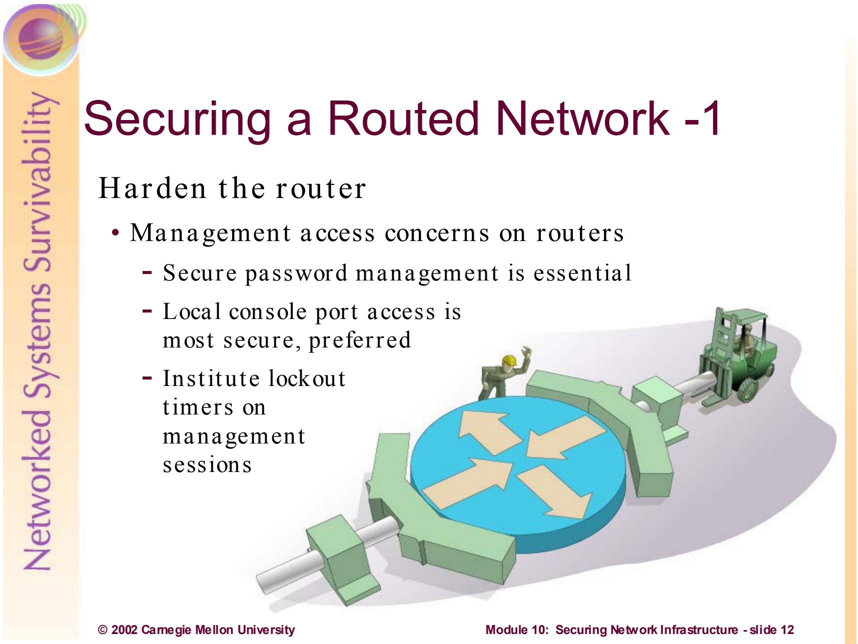

Networked Systems Survivability

Securing a Routed Network -1

Harden the router

- Management access concerns on routers
  - Secure password management is essential
  - Local console port access is most secure, preferred
  - Institute lockout timers on management sessions

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 12

Routers typically are special purpose network infrastructure devices that direct network data messages, or packets, based on internal addresses and tables of routes, or known destinations that serve certain addresses. Directing data between networks or portions of a network is the primary purpose of a router. Routers can also be implemented on standard PCs that have at least 2 network cards and special software. Windows 2000 comes with router software built in--as does almost all versions of Linux. Our discussion will focus on Cisco routers—as they are by far the most common and deployed routers in the world.

Effective password management is critical throughout your network, but particularly important for securing routers and other infrastructure devices.

As a general best practice, configure durable (difficult to guess or crack) passwords on the router or the firewall--if the router is being used as one. For each password use the following guidelines: be at least eight characters long, not be words, not begin with a number, and include at least one character from the sets of letters, numbers and all other characters (e.g., ,./<>;’:”[]\{}|~!@#$%^&*()_+`-=). Consider using different passwords for each router and each firewall. Change passwords at least once every 90 days [NSA2].

There are two password protection schemes in Cisco IOS. Type 7 uses the Cisco-defined encryption algorithm which is known to the commercial security community to be weak. Type 5 uses an MD5 hash which is much stronger. Cisco recommends that Type 5 encryption be used instead of Type 7 where possible (see Configuring Passwords and Privileges[14] in the Cisco IOS Security Configuration Guide[15]). Type 7 encryption is used by the **enable password**, **username**, and line **password** commands. You should know that weak Type 7 passwords can be cracked easily—even on your Palm compatible PDA![16] Type 5 is for enable secret passwords. It is advisable to always use enable secret passwords when configuring access to enable mode (a.k.a. privileged Exec)--leave standard enable passwords unset. This protects the privileged EXEC mode which is similar to administrator or root on Windows and Unix hosts.

There are two types of management access: local and remote. Local access usually involves a direct connection to a console port on the router with a dumb terminal or a laptop computer. Remote access typically involves allowing telnet, SNMP, or even http connections to the router from some computer on

[14] http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt5/scdpass.htm
[15] http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/
[16] http://www.atstake.com/research/tools/cisco.zip

the same subnet or a different subnet (it is not uncommon to remotely access routers from the other side of the country or even the globe). It is recommended to only allow local access because during remote access all telnet passwords or SNMP community strings are sent in the clear to the router. If an attacker can collect network traffic during remote access then he can capture passwords or community strings [NSA1].

The console (con) port is the default location for performing router management and configuration. It is okay to leave a connection to the console port attached all the time, but that terminal (or computer) should be standalone, and protected from unauthorized access. The connection to the console port should not be left logged in. Configure the console line to time out, so that if an administrator forgets to log out, the router will log him or her out automatically.

Following are example password and time out configurations for Cisco routers [NSA2]:

Configure the console line and the virtual terminal lines (remote telnet) on the router to time out a session (after 5 minutes) and to require a password at login.

```
Router(config)# line con 0
Router(config-line)# exec-timeout 5 0
Router(config-line)# login
Router(config)# line vty 0 4
Router(config-line)# exec-timeout 5 0
Router(config-line)# login
```

Configure the Enable Secret password, which is protected with an MD5-based algorithm.

```
Router(config)# enable secret 0 2manyRt3s
```

Configure passwords for the console line, the auxiliary line and the virtual terminal lines. Use a different password for the console line versus the virtual terminal lines.

```
Router(config)# line con 0
Router(config-line)# password Soda-4-jimmY
Router(config)# line vty 0 4
Router(config-line)# password Dots-4-georg3
```

Provide a basic protection for the line passwords by using the following global configuration command. This will keep passersby from reading your passwords when they are displayed on the screen

```
Router(config)# service password-encryption
```

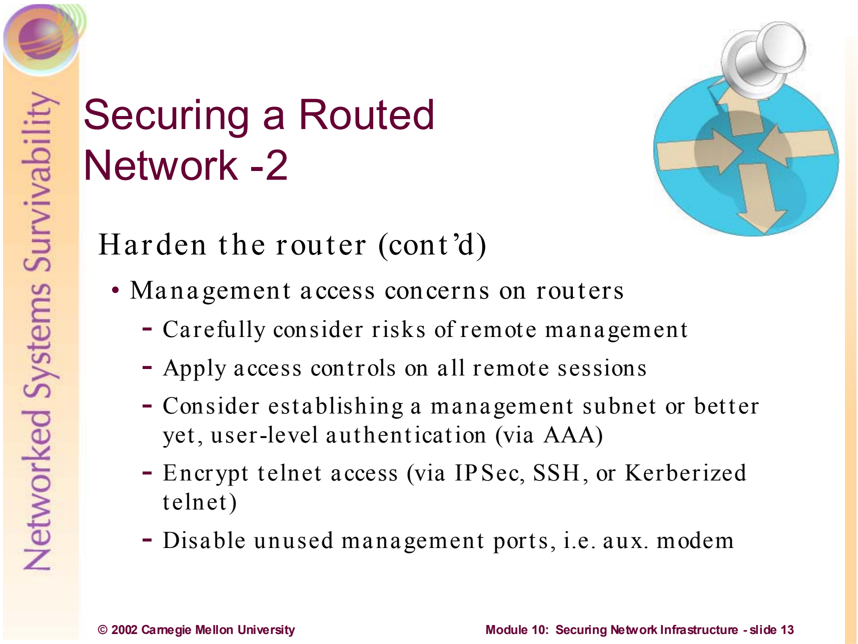

Networked Systems Survivability

Securing a Routed Network -2

Harden the router (cont'd)

- Management access concerns on routers
  - Carefully consider risks of remote management
  - Apply access controls on all remote sessions
  - Consider establishing a management subnet or better yet, user-level authentication (via AAA)
  - Encrypt telnet access (via IPSec, SSH, or Kerberized telnet)
  - Disable unused management ports, i.e. aux. modem

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 13

If remote administration is required, it is important to mitigate the inherent risks to overall network security. Cisco is supporting (and recommending the use of) Authentication, Authorization, and Accounting (AAA) which is an emerging IETF security specification.[17] It is designed for controlling access, privileges, and logging of user activities on a router. Authentication is the mechanism for identifying users before allowing access to a network component. Authorization is the method used to describe what a user has the right to do once he has authenticated to the router. Accounting is the component that allows for logging and tracking of user and traffic activities on the router which can be used later for resource tracking or trouble shooting. AAA requires significant architectural design considerations to implement, however it is certainly superior to traditional methods of router management access.[18]

If you must use traditional virtual terminal telnet remote access, you should set up network access controls on the router to restrict the traffic that can access the router's vty ports. A good practice is to set up one--or just a few--administration hosts (with static IP addresses) that are part of a separate management subnet. Then set access controls on the router to allow telnet connections only from these hosts.

Remember your telnet sessions are susceptible to sniffing because telnet is "pass in the clear." Because of this, you should institute some method of encrypting the traffic as it passes through the network. There are many technologies (see module 12) for accomplishing this:

- **IPSec**: because of its interoperability and robustness, may be the preferred means of encrypting management sessions on your routers. See section 5.2 in the NSA's Router Security Configuration guide for step by step instructions on how to setup IPSec sessions between Windows 2000 clients and Cisco Routers (complete with screen shots!)[19]
- **Secure Shell (SSH):** is widely used for encrypting remote sessions. Has proven security track record (especially in the Unix community), however it has only recently become widely supported by Cisco. Also, some serious vulnerabilities exists with certain versions of the SSH protocol (and

[17] http://www.ietf.org/html.charters/aaa-charter.html
[18] http://www.cisco.com/univercd/cc/td/doc/product/software/ios113ed/113ed_cr/secur_c/scprt1/
[19] http://nsa1.www.conxion.com/cisco/guides/cis-2.pdf

especially) as it's implemented on Cisco equipment.[20] If you choose to use SSH, make sure your equipment has been patched.

- **Kerberized Telnet:**[21] is an authentication and encryption methodology that is less widely used than the above methods, however it should be considered because of the proven security benefits of Kerberos—covered later in this module. Cisco supports this method (as part of AAA) and there are many client programs available for free download.[22]

All unused management ports should be disabled. In particular, the aux (modem) port should be disabled. Only if absolutely required should a modem be connected to the aux port as a backup or remote access method to the router. Attackers using simple war-dialing software will eventually find the modem, so it is necessary to apply access controls to the aux port—if you must utilize it. For better security, IOS callback features should be used [NSA1].

[20] http://www.cisco.com/warp/public/707/SSH-multiple-pub.html
[21] http://web.mit.edu/is/help/ktelnet/
[22] http://www.stacken.kth.se/~thn/ktelnet/

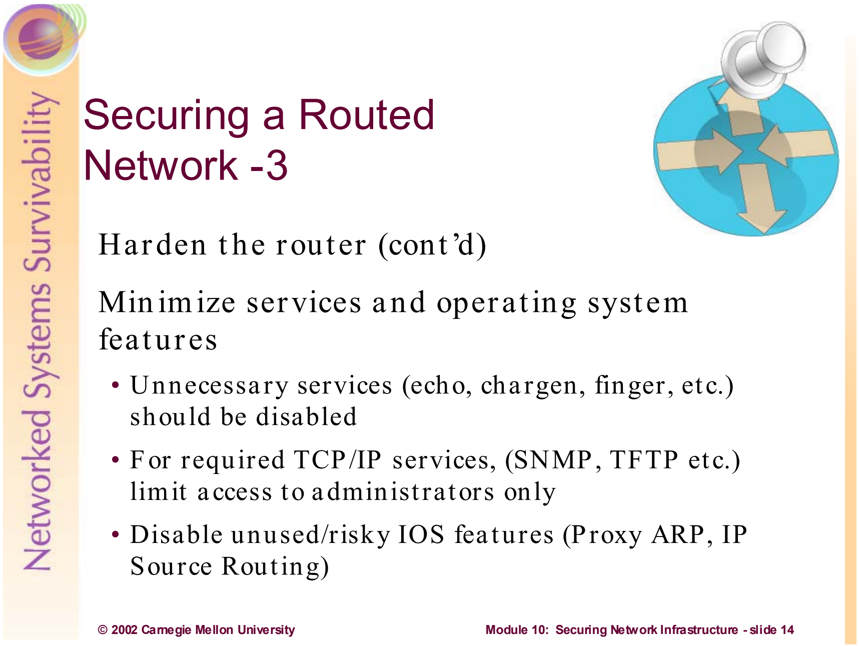

Networked Systems Survivability

Securing a Routed Network -3

Harden the router (cont'd)

Minimize services and operating system features

- Unnecessary services (echo, chargen, finger, etc.) should be disabled
- For required TCP/IP services, (SNMP, TFTP etc.) limit access to administrators only
- Disable unused/risky IOS features (Proxy ARP, IP Source Routing)

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 14

It is important to shut down all unneeded TCP/UDP services on the router. Services that are not running can not cause problems, or be used as the basis for attacks. Also, you will be freeing up memory and processing cycles by minimizing services.

You may require some TCP/IP services (like TFTP or SNMP) as part of your network management and administrative tasks. Use these services with caution—they can open the door to intruders if they are not tightly controlled.

There are many features in Cisco IOS that are enabled by default (for legacy reasons) however, they present security risks and should be disabled—see below.

The show processes command can help to show active information about the servers on the router. The following commands show how to disable the following servers: TCP/UDP small servers (echo, discard, daytime, chargen), bootps, finger, http, identd and snmp.

```
Router(config)# no service tcp-small-servers
Router(config)# no service udp-small-servers
Router(config)# no ip bootp server
Router(config)# no service finger
Router(config)# no ip http server
Router(config)# no ip identd
Router(config)# no snmp-server community <community string>
```

If SNMP on the router is required, use the following commands to clear out any SNMP servers with default community strings.

```
Router(config)# no snmp-server community public
Router(config)# no snmp-server community private
```

Then set up the SNMP server with a community string that is difficult to guess. Also, if possible, allow only read-only access to the server; do not allow read-write access to the server. Apply an access-list to the server. Refer to the following section on TCP/IP Filters for discussion of an access-list for SNMP in more detail. The following command is an example.

```
Router(config)# snmp-server community S3cr3t-str1n9 ro 10
```

The following commands disable the following (unneeded) services: Cisco Discovery Protocol (CDP), remote configuration downloading, source routing and zero subnet.

```
Router(config)# no cdp run
Router(config)# no service config
Router(config)# no ip source-route
Router(config)# no ip subnet-zero
```

Following is a table of Risky Services and IOS Features [NSA 1]:

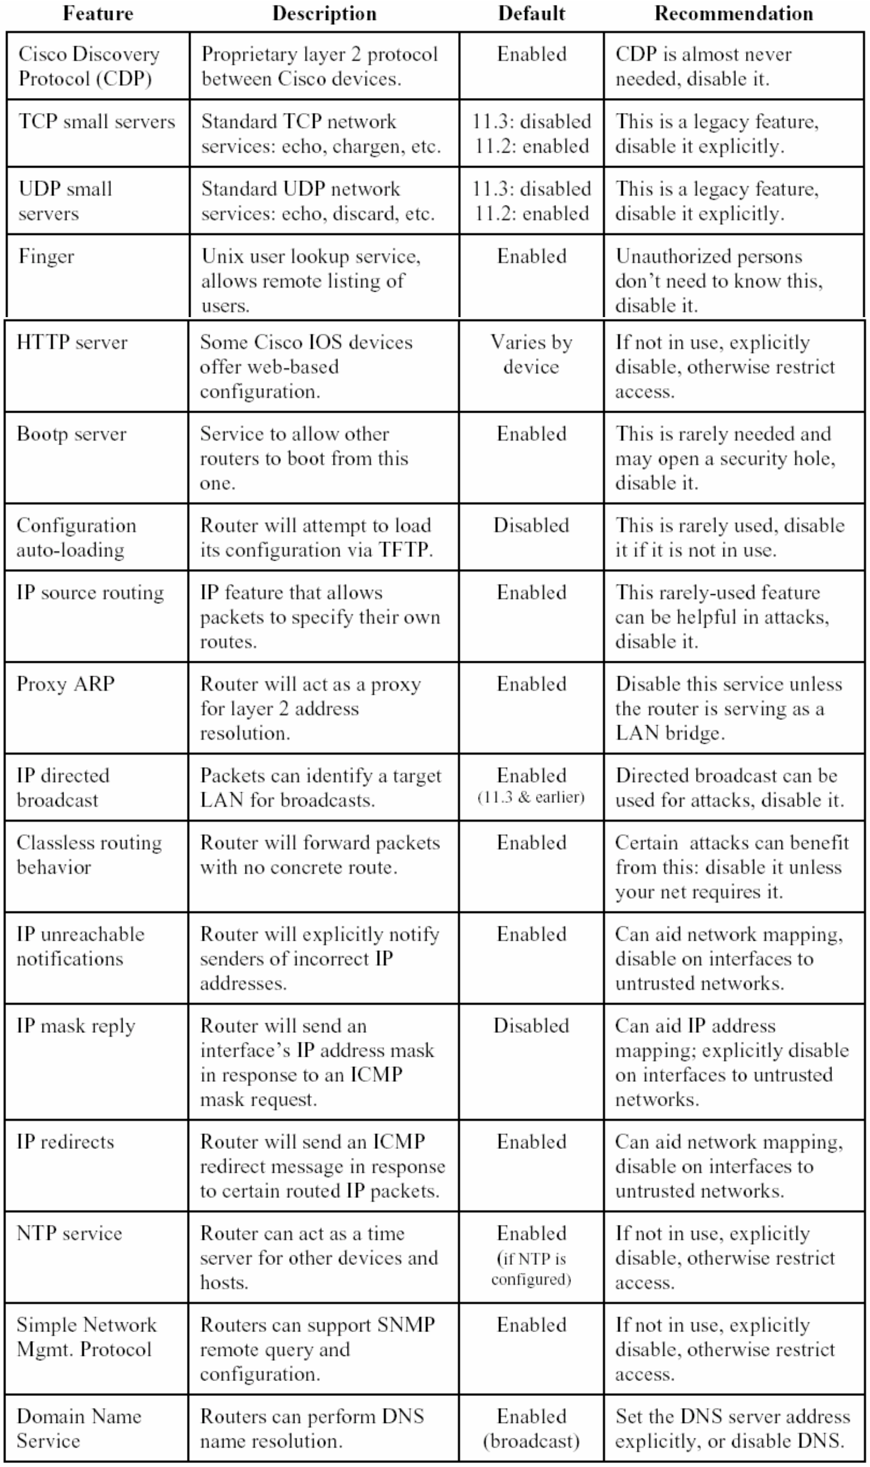

Feature	Description	Default	Recommendation
Cisco Discovery Protocol (CDP)	Proprietary layer 2 protocol between Cisco devices.	Enabled	CDP is almost never needed, disable it.
TCP small servers	Standard TCP network services: echo, chargen, etc.	11.3: disabled 11.2: enabled	This is a legacy feature, disable it explicitly.
UDP small servers	Standard UDP network services: echo, discard, etc.	11.3: disabled 11.2: enabled	This is a legacy feature, disable it explicitly.
Finger	Unix user lookup service, allows remote listing of users.	Enabled	Unauthorized persons don't need to know this, disable it.
HTTP server	Some Cisco IOS devices offer web-based configuration.	Varies by device	If not in use, explicitly disable, otherwise restrict access.
Bootp server	Service to allow other routers to boot from this one.	Enabled	This is rarely needed and may open a security hole, disable it.
Configuration auto-loading	Router will attempt to load its configuration via TFTP.	Disabled	This is rarely used, disable it if it is not in use.
IP source routing	IP feature that allows packets to specify their own routes.	Enabled	This rarely-used feature can be helpful in attacks, disable it.
Proxy ARP	Router will act as a proxy for layer 2 address resolution.	Enabled	Disable this service unless the router is serving as a LAN bridge.
IP directed broadcast	Packets can identify a target LAN for broadcasts.	Enabled (11.3 & earlier)	Directed broadcast can be used for attacks, disable it.
Classless routing behavior	Router will forward packets with no concrete route.	Enabled	Certain attacks can benefit from this: disable it unless your net requires it.
IP unreachable notifications	Router will explicitly notify senders of incorrect IP addresses.	Enabled	Can aid network mapping, disable on interfaces to untrusted networks.
IP mask reply	Router will send an interface's IP address mask in response to an ICMP mask request.	Disabled	Can aid IP address mapping; explicitly disable on interfaces to untrusted networks.
IP redirects	Router will send an ICMP redirect message in response to certain routed IP packets.	Enabled	Can aid network mapping, disable on interfaces to untrusted networks.
NTP service	Router can act as a time server for other devices and hosts.	Enabled (if NTP is configured)	If not in use, explicitly disable, otherwise restrict access.
Simple Network Mgmt. Protocol	Routers can support SNMP remote query and configuration.	Enabled	If not in use, explicitly disable, otherwise restrict access.
Domain Name Service	Routers can perform DNS name resolution.	Enabled (broadcast)	Set the DNS server address explicitly, or disable DNS.

Care should be taken to review all of these features and services so that your routers will be as secure as is possible.

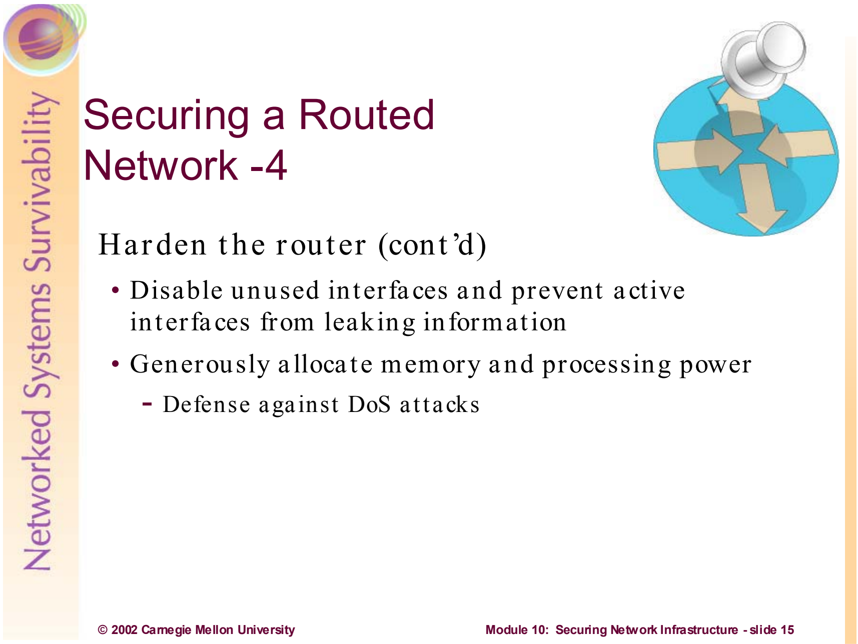

Networked Systems Survivability

Securing a Routed Network -4

Harden the router (cont'd)

- Disable unused interfaces and prevent active interfaces from leaking information
- Generously allocate memory and processing power
  - Defense against DoS attacks

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 15

It is considered a security best practice to disable unused LAN and WAN interfaces on your routers. Also, care should be given to harden active interfaces. Below are some commands that are examples for doing so:

The following command disables a router interface.

```
Router(config-if)# shutdown
```

Secure each and every active interface on the router from Smurf attacks, ad-hoc routing, and access-list queries with the following commands.

```
Router(config-if)# no ip directed-broadcast
Router(config-if)# no ip proxy-arp
Router(config-if)# no ip unreachables
```

On today's Internet, Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks are an every day occurrence. The most common variant of these are packet flooding attacks. In this case, packets are deluged on a target causing it to be unable to respond to legitimate traffic. In February of 2000, large Internet web sites like Yahoo.com, Ebay.com, Amazon.com, Buy.com, and CNN.com were all victims of major DDoS attacks.[23] The loss suffered from these attacks was estimated at more than $1 billion[24]. In response, many of these firms (and others as well) invested heavily in infrastructure and bandwidth. The idea being to overcompensate on all aspects of network infrastructure—thereby defeating DDoS and other flooding attacks by simply requiring more traffic then the attackers can feasibly generate (to be noticed as a problem). Overcompensating memory and processing power on routers is recommended—if the mission of the organization makes its availability critical enough to justify the expense.

[23] http://www.computerworld.com/cwi/story/0,1199,NAV47_STO43010,00.html

[24] http://dgl.com/itinfo/2000/it000214.html

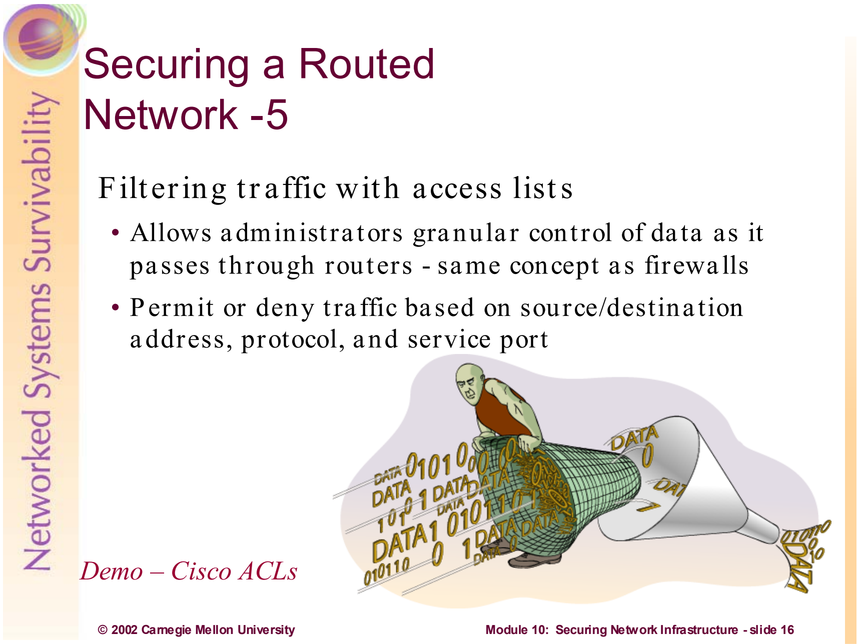

Networked Systems Survivability

Securing a Routed Network -5

Filtering traffic with access lists

- Allows administrators granular control of data as it passes through routers - same concept as firewalls
- Permit or deny traffic based on source/destination address, protocol, and service port

Demo – Cisco ACLs

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 16

A packet filter for TCP/IP services provides control of the data transfer between networks based on addresses and protocols. Routers can apply filters in different ways.

Some routers have filters that apply to network services in both inbound and outbound directions, while others have filters that apply only in one direction. (Many services are bi-directional. For example, a user on System A telnets to System B, and System B sends some type of response back to System A. So, some routers need two filters to handle bi-directional services.) Most routers can filter on one or more of the following: source IP address, source port, destination IP address, destination port, and protocol type. Some routers can even filter on any bit or any pattern of bits in the IP header. However, routers do not have the capability to filter on the content of services (e.g. FTP file name).

Packet filters are especially important for routers that act as the gateway between trusted and untrusted networks. In that role, the router can enforce security policy, rejecting protocols and restricting ports according to the policies of the trusted network. Filters are also important for their ability to enforce addressing constraints. For example, a router should enforce the constraint that packets sent from the Firewall or protected network out to the Internet must bear a source address within a particular range. This is sometimes called *egress filtering*. Similarly, the router should enforce the constraint that packets arriving from the Internet must bear a source address outside the range valid for the protected network. This is called *ingress filtering*.

Two key characteristics of TCP/IP packet filters are length and ordering. A filter consists of one or more rules, with each rule either accepting or denying a certain set of packets. The number of rules in a filter determines its length. Generally, as the length grows the filter becomes more complex and more difficult to troubleshoot. The order of the rules in a packet filter is critical. When the router analyzes a packet against a filter, the packet is compared to each filter rule in sequential order. If a match is found then the packet is either permitted or denied and the rest of the filter is ignored. If no match is found then the packet is denied due to the implicit deny rule at the end of the filter. You must carefully create filter rules in the proper order so that all packets are treated according to the intended security policy. One method of ordering involves placing those rules that will handle the bulk of the traffic as close to the beginning of the filter as possible. Consequently, the length and ordering of a packet filter rule set can affect the router's performance.

On Cisco routers, Access lists can be created for many different kinds of protocols. For IP traffic, there are two types of access lists available: standard and extended. Standard access lists only allow source IP address filtering. Extended access lists can permit or deny packets based on their protocols, source or destination IP addresses, source or destination TCP/UDP ports, or ICMP or IGMP message types. Extended access lists also support selective logging. Both standard and extended IP access lists can be applied to router interfaces, vty lines (for remote access), IPSec, routing protocols, and many router features. Only standard IP access lists can be applied to SNMP.

The basic structure for an access list rule is shown below.

access-list *list-number* **{deny** | **permi**t} *condition*

The access list number tells Cisco IOS which access list the rule should be a part of, and what kind of access list it is. The condition field, which is different for each kind of access list, specifies which packets match the rule. Conditions typically involve protocol information and addresses, but do not involve application-level information. The following is the syntax for a statement (rule) in a standard IP access list:

access-list *list-number* **{deny** | **permi**t} *source [source-wildcar*d] **[log**]

where *list-number* is the number of the access list and can be any decimal number from 1 to 99.

deny denies access if the condition is matched.

permit permits access if the condition is matched.

source is the IP address of the network or host from which the packet is being sent.

source-wildcard is the wildcard bits to be applied to the *sourc*e.

log, if present, causes a message about the packet that matches the statement to be logged

Extended IP access list provide more granularity and control.[25] The following is simplified syntax for a statement in an extended IP access list:

access-list *list-number* **{deny** | **permi**t} *protocol source source-wildcard source-qualifiers*

destination destination-wildcard destination-qualifiers [**log** | **log-inpu**t][26] [NSA 1]

[25] http://www.networkcomputing.com/907/907ws1.html

[26] http://img.cmpnet.com/nc/907/graphics/access.pdf

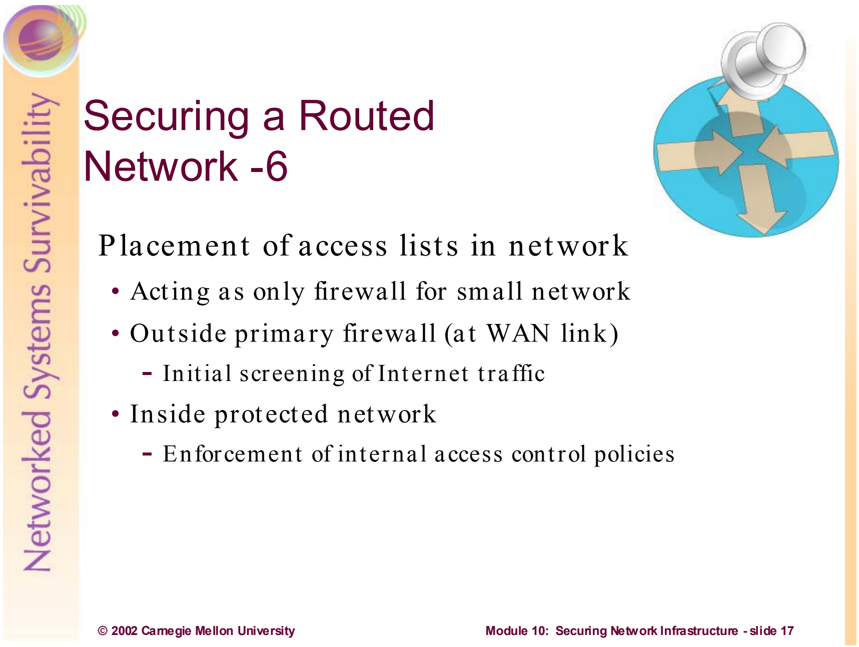

Networked Systems Survivability

Securing a Routed Network -6

Placement of access lists in network

- Acting as only firewall for small network
- Outside primary firewall (at WAN link)
  - Initial screening of Internet traffic
- Inside protected network
  - Enforcement of internal access control policies

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 17

A router provides a capability to help secure the perimeter of a protected network. It can do this by itself. The diagram at right shows a typical topology with the router being the component that connects the protected network to the Internet.

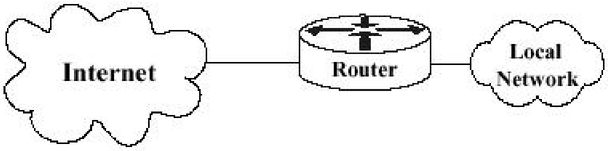

A router can also be used as part of defense-in-depth approach as shown in the diagram below. It acts as the first line of defense and is known as a screening router. It contains a static route that passes all connections intended for the protected network to the firewall. The firewall provides additional access control over the content of the connections. It can also perform user authentication. This approach is recommended over using only a router because it offers more security.

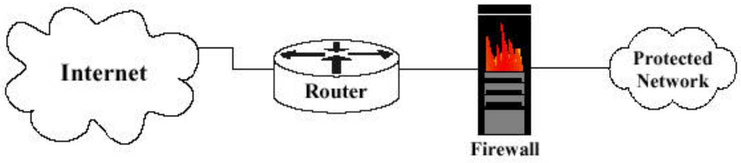

Another approach is to position one router at the connection between the local premises and the Internet, and then another router between the firewall and the protected network. This configuration offers two points at which policy can be enforced. It also offers an intermediate area, often called the de-militarized zone (DMZ) between the two routers. The DMZ is often used for servers that must be accessible from the Internet or other external network [NSA1].

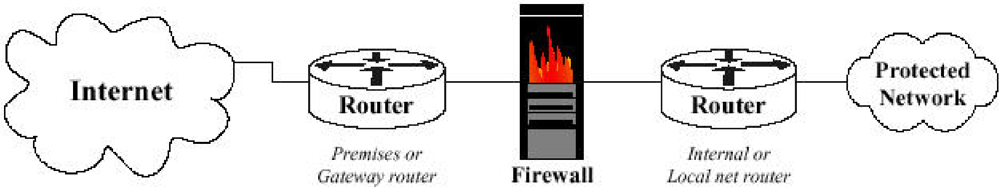

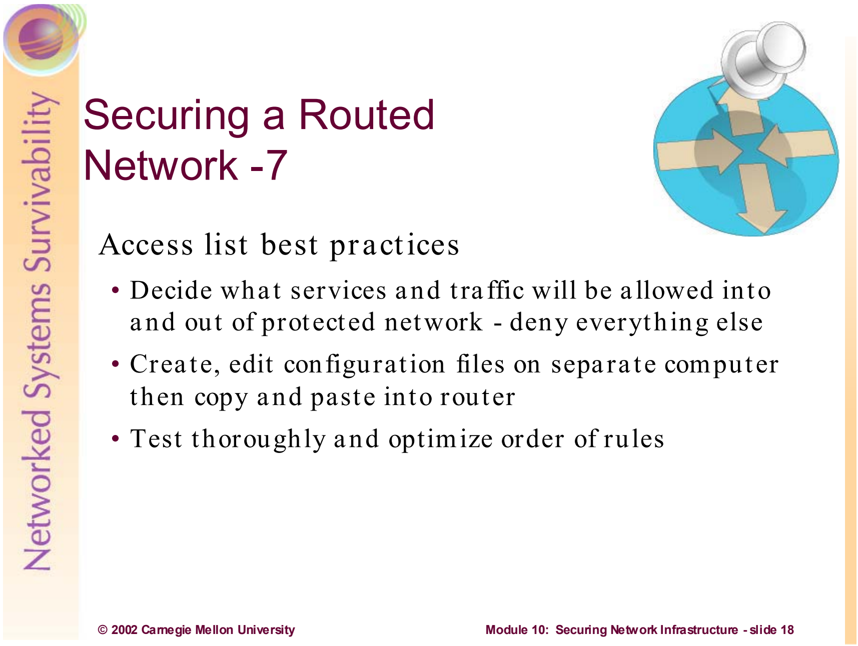

Networked Systems Survivability

Securing a Routed Network -7

Access list best practices

- Decide what services and traffic will be allowed into and out of protected network - deny everything else
- Create, edit configuration files on separate computer then copy and paste into router
- Test thoroughly and optimize order of rules

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 18

Access list are (or should be) just a technological enforcement of written security policies. These policies should specify what traffic should be allowed into and out of the protected network. Additionally, policies should define what traffic should be allowed in and out of internal subnets.

Due to limited editing capability on the Cisco router, you cannot easily modify access lists. Thus, whenever you need to change an access list, it is best to build it offline on a separate computer. When the access list is ready you can cut and paste the access list via a connection to the router. Since the original access list is still on the router, you must purge it before adding the updated access list. Remember to apply the new access list to the appropriate interface—it won't be active until you do so.

It is always best to test Access lists carefully before deploying them on production networks. This is not always possible however, and steps should be taken to facilitate the speedy reversion back to the old configuration.

Administrators should review the router's logs and counters to determine which access list rules are being tested most frequently. Once this has been determined, these rules should be moved to the top of the access list. Doing so will decrease processing cycles on the router. Be very careful when doing this operation—changing the order of access lists can significantly impact the desired result.

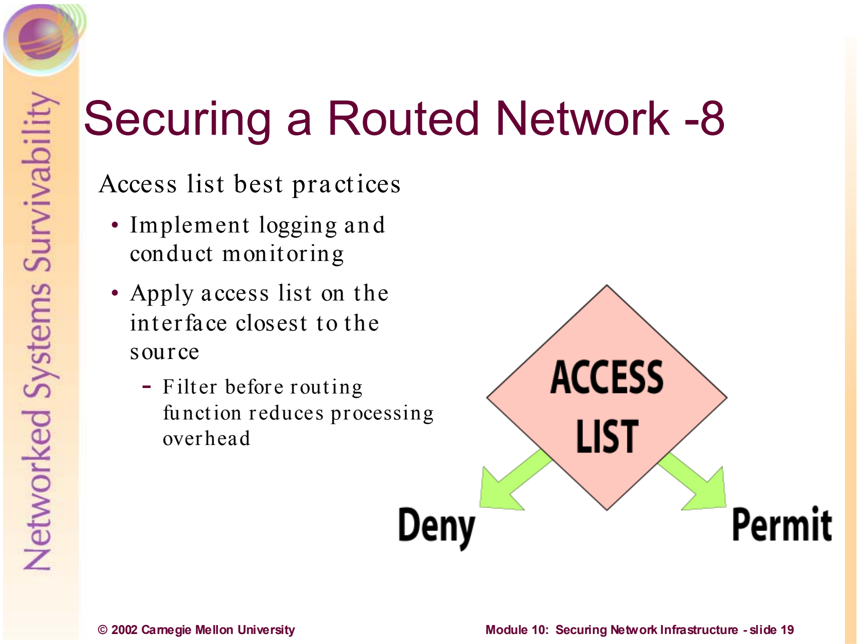

Networked Systems Survivability

Securing a Routed Network -8

Access list best practices

- Implement logging and conduct monitoring
- Apply access list on the interface closest to the source
  - Filter before routing function reduces processing overhead

ACCESS LIST
Deny
Permit

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 19

Logging on a router or a firewall offers several benefits. It informs the administrator if the router or the firewall is working properly or has been compromised. It can also show what types of attacks are being attempted against the router, the firewall or the protected network. The following are recommendations for logging and debugging:

- Send the most serious level of logs to the console on the router or the firewall in order to alert the administrator.
- Send the logs to a log host, which should be a dedicated computer on the protected network whose only job is to receive logs. The log host should have all unnecessary servers and accounts disabled except for syslog.
- Configure the router or the firewall to include more specific time information in the logging and in the debugging. Direct the router or the firewall to at least two different, reliable network time protocol (NTP) servers to ensure accuracy and availability of time information. Set all NTP messages with the same IP source address of an interface on the internal network. This configuration will allow the administrator to create a TCP/IP filter that allows time information only from the internal IP address of the router or the firewall to the external NTP servers. This filter will help to prevent spoofing or flooding NTP messages to the router or the firewall. Include a more specific timestamp in each log message and each debug message. This will allow an administrator to trace network attacks more credibly.
- By default, a log message contains the IP address of the interface it uses to leave the router or the firewall. Instead, set all log messages with the same IP source address of an interface on the internal network, regardless of which interface the messages use. This configuration will allow the administrator to create a TCP/IP filter that allows logs only from the internal IP address of the router or the firewall to the logging host. This filter will help to prevent spoofing or flooding log messages to the logging host.
- Finally, consider also sending the logs to a dedicated printer to deal with worst-case scenarios, e.g., failure of the log host [NSA2].

The diagram below shows how access lists work when applied to router interfaces, using the router East as an example [NSA1].

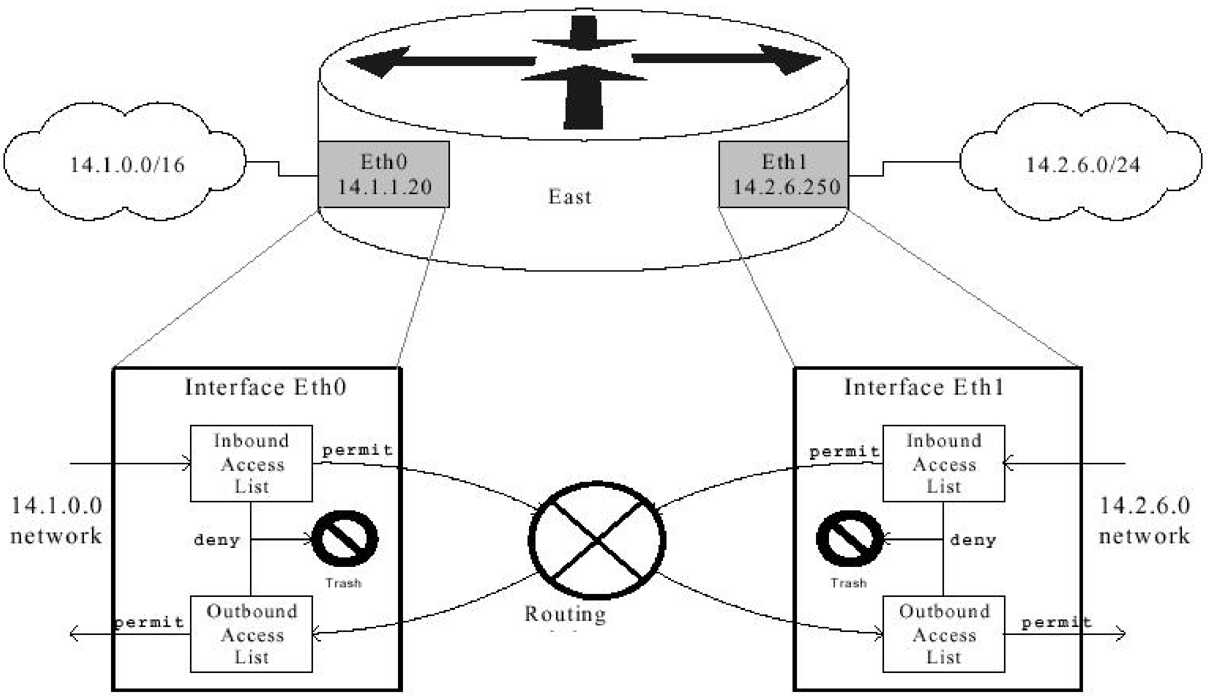

The diagram also illustrates why access lists should be applied closest to the source of the traffic you are filtering. Inbound access lists are more efficient and are preferred over outbound access lists. This is because packets that are denied are trashed before any routing has occurred, thereby freeing up memory and processing cycles on the router. In some cases, outbound access lists are unavoidable; however they should be minimized wherever possible.

Enter the following command to enable the logging capability on a Cisco router:

```
Router(config)# logging on
```

*** In each access list there must be at least one **permit** statement. Otherwise, an access list with no **permit** statements will block all network traffic wherever it is applied*** [NSA2].

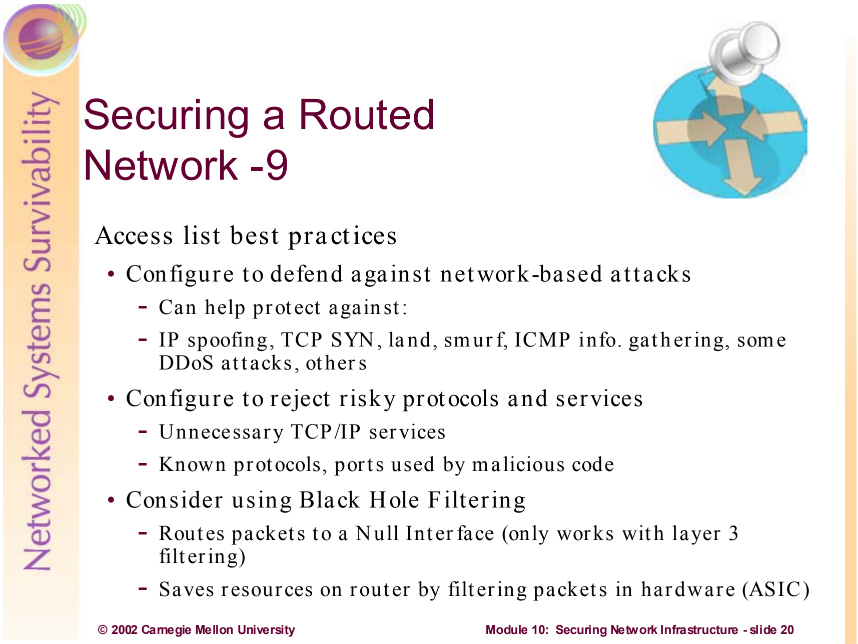
Networked Systems Survivability

Securing a Routed Network -9

Access list best practices

- Configure to defend against network-based attacks
  - Can help protect against:
  - IP spoofing, TCP SYN, land, smurf, ICMP info. gathering, some DDoS attacks, others
- Configure to reject risky protocols and services
  - Unnecessary TCP/IP services
  - Known protocols, ports used by malicious code
- Consider using Black Hole Filtering
  - Routes packets to a Null Interface (only works with layer 3 filtering)
  - Saves resources on router by filtering packets in hardware (ASIC)

© 2002 Carnegie Mellon University   Module 10: Securing Network Infrastructure - slide 20

Access lists can be helpful in defending against network based attacks.

- Provide IP address spoof protection for the protected network. For inbound traffic do not allow any IP packet that contains an IP address in the source IP address field from the following: the protected network, any local host address (127.0.0.0 –127.255.255.255), any reserved address (10.0.0.0 – 10.255.255.255, 172.16.0.0 –172.31.255.255, 192.168.0.0 – 192.168.255.255), or any multicast address (224.0.0.0 – 239.255.255.255). For outbound traffic allow IP traffic from the protected network and do not allow IP traffic that contains an external IP address in the source IP address field.
- Protect the router or the firewall from the Land Attack. This attack involves sending a packet to the router with the same IP address in the source address and destination address fields and with the same port number in the source port and destination port fields. This attack can cause a denial of service.
- Protect the router or the firewall from the TCP SYN Attack. The TCP SYN Attack involves transmitting a volume of connections that cannot be completed at the destination. This attack causes the connection queues on the router or the firewall to fill up, thereby denying service to legitimate TCP traffic.
- Protect the router, the firewall or the protected network from unnecessary ICMP traffic. There are a variety of ICMP message types, and some are associated with programs. Some message types are used for network management and are automatically generated and interpreted by network devices. For example, the ping program works with message type Echo. With Echo packets an attacker can create a map of the protected networks behind the router or the firewall. Also, he can perform a denial of service attack by flooding the router, the firewall or the hosts on the protected network with Echo packets. With Redirect packets the attacker can cause changes to a host’s routing tables.
- For outbound ICMP traffic, one should allow the message types Echo, Parameter Problem and Source Quench. Otherwise, block all other ICMP message types going outbound. With Echo packets users will be able to ping external hosts. Parameter Problem packets and Source Quench packets improve connections by informing about problems with packet headers and by slowing down traffic when it is necessary. For inbound ICMP traffic, one should allow the following message types: Echo Reply, Destination Unreachable, Source Quench, Time Exceeded and Parameter Problem. Otherwise, block all other ICMP message types coming inbound.

- Protect the router, the firewall or the protected network from inbound traceroute. Traceroute is a utility that prints the IP addresses of the routers that handle a packet as the packet hops along the network from source to destination. On Unix operating systems traceroute uses UDP packets and causes routers along the path to generate ICMP message types Time Exceeded and Unreachable. Similar to ICMP Echo packets, an attacker can use traceroute to create a map of the protected network behind the router or the firewall. [NSA2]

Cisco IOS example configurations for using access lists to defend against some attacks [NSA2]:

The following commands show how to protect the router from the Land Attack:

```
Router(config)# access-list 101 deny ip host 198.26.171.178 host 198.26.171.178 log
Router(config)# access-list 101 permit ip any any
Router(config)# interface serial2/1
Router(config-if)# description "external interface"
Router(config-if)# ip address 198.26.171.178 255.255.255.248
Router(config-if)# ip access-group 101 in
```

Protect the router against the TCP SYN Attack for the following two scenarios: blocking external access and limited external access. Below is an example for blocking external access on a Cisco router. The access list blocks packets from any external network that have only the SYN flag set. Thus, it allows traffic from TCP connections that were established from the protected network (e.g., 14.2.6.0), and it denies anyone coming from any external network from starting any TCP connection.

```
Router(config)# access-list 100 permit tcp any 14.2.6.0 0.0.0.255
established
Router(config)# access-list 100 deny ip any any log
Router(config)# interface serial0/0
Router(config-if)# description "external interface"
Router(config-if)# ip access-group 100 in
```

Use the TCP intercept feature to provide additional protection from SYN attacks. This feature allows only reachable external hosts to initiate connections to a host on the internal network. In intercept mode, the router intercepts each TCP connection establishment and determines if the address from which the connection is being initiated is reachable. If the host is reachable, the router allows the connection to be established; otherwise, it prevents the connection. Because spoofed IP packets are often part of SYN attacks, TCP intercept can be an effective mechanism of keeping these bogus (forged by malicious software) packets out of your protected network. [NSA1].

Many of these are known attack programs that can cause damage if allowed inside your protected network. Others, like NetBIOS ports 137 and 138, (used in Windows networks for resource sharing) are notorious for leaking information about your network to intruders.

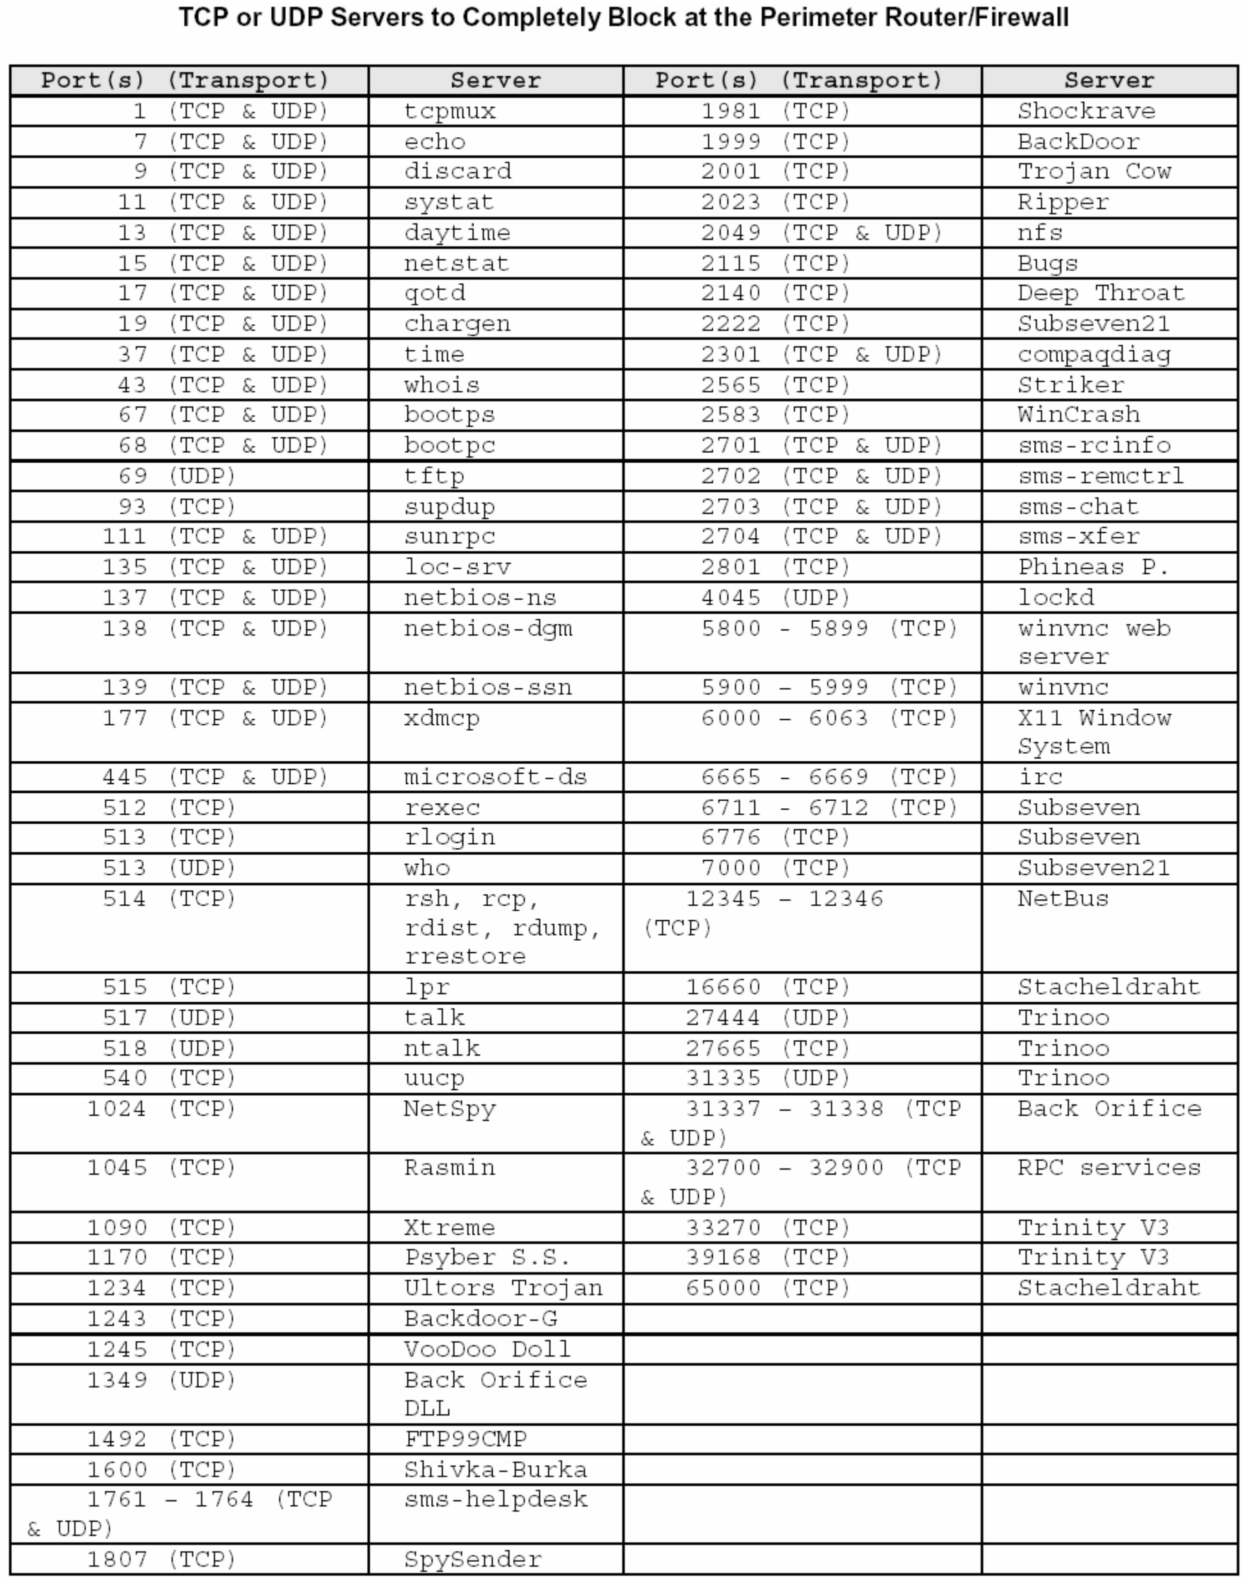

TCP or UDP Servers to Completely Block at the Perimeter Router/Firewall

Port(s) (Transport)	Server	Port(s) (Transport)	Server
1 (TCP & UDP)	tcpmux	1981 (TCP)	Shockrave
7 (TCP & UDP)	echo	1999 (TCP)	BackDoor
9 (TCP & UDP)	discard	2001 (TCP)	Trojan Cow
11 (TCP & UDP)	systat	2023 (TCP)	Ripper
13 (TCP & UDP)	daytime	2049 (TCP & UDP)	nfs
15 (TCP & UDP)	netstat	2115 (TCP)	Bugs
17 (TCP & UDP)	qotd	2140 (TCP)	Deep Throat
19 (TCP & UDP)	chargen	2222 (TCP)	Subseven21
37 (TCP & UDP)	time	2301 (TCP & UDP)	compaqdiag
43 (TCP & UDP)	whois	2565 (TCP)	Striker
67 (TCP & UDP)	bootps	2583 (TCP)	WinCrash
68 (TCP & UDP)	bootpc	2701 (TCP & UDP)	sms-rcinfo
69 (UDP)	tftp	2702 (TCP & UDP)	sms-remctrl
93 (TCP)	supdup	2703 (TCP & UDP)	sms-chat
111 (TCP & UDP)	sunrpc	2704 (TCP & UDP)	sms-xfer
135 (TCP & UDP)	loc-srv	2801 (TCP)	Phineas P.
137 (TCP & UDP)	netbios-ns	4045 (UDP)	lockd
138 (TCP & UDP)	netbios-dgm	5800 - 5899 (TCP)	winvnc web server
139 (TCP & UDP)	netbios-ssn	5900 – 5999 (TCP)	winvnc
177 (TCP & UDP)	xdmcp	6000 – 6063 (TCP)	X11 Window System
445 (TCP & UDP)	microsoft-ds	6665 - 6669 (TCP)	irc
512 (TCP)	rexec	6711 - 6712 (TCP)	Subseven
513 (TCP)	rlogin	6776 (TCP)	Subseven
513 (UDP)	who	7000 (TCP)	Subseven21
514 (TCP)	rsh, rcp, rdist, rdump, rrestore	12345 – 12346 (TCP)	NetBus
515 (TCP)	lpr	16660 (TCP)	Stacheldraht
517 (UDP)	talk	27444 (UDP)	Trinoo
518 (UDP)	ntalk	27665 (TCP)	Trinoo
540 (TCP)	uucp	31335 (UDP)	Trinoo
1024 (TCP)	NetSpy	31337 – 31338 (TCP & UDP)	Back Orifice
1045 (TCP)	Rasmin	32700 – 32900 (TCP & UDP)	RPC services
1090 (TCP)	Xtreme	33270 (TCP)	Trinity V3
1170 (TCP)	Psyber S.S.	39168 (TCP)	Trinity V3
1234 (TCP)	Ultors Trojan	65000 (TCP)	Stacheldraht
1243 (TCP)	Backdoor-G		
1245 (TCP)	VooDoo Doll		
1349 (UDP)	Back Orifice DLL		
1492 (TCP)	FTP99CMP		
1600 (TCP)	Shivka-Burka		
1761 – 1764 (TCP & UDP)	sms-helpdesk		
1807 (TCP)	SpySender		

Here are some other recommendations for blocking/limiting network traffic: [NSA2]

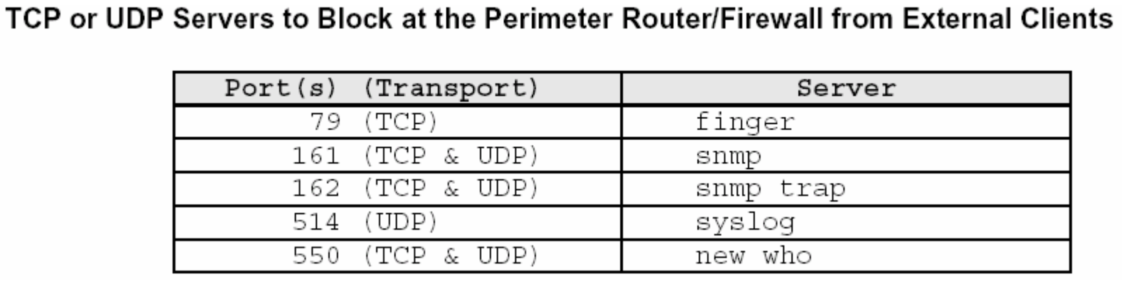

TCP or UDP Servers to Block at the Perimeter Router/Firewall from External Clients

Port(s) (Transport)	Server
79 (TCP)	finger
161 (TCP & UDP)	snmp
162 (TCP & UDP)	snmp trap
514 (UDP)	syslog
550 (TCP & UDP)	new who

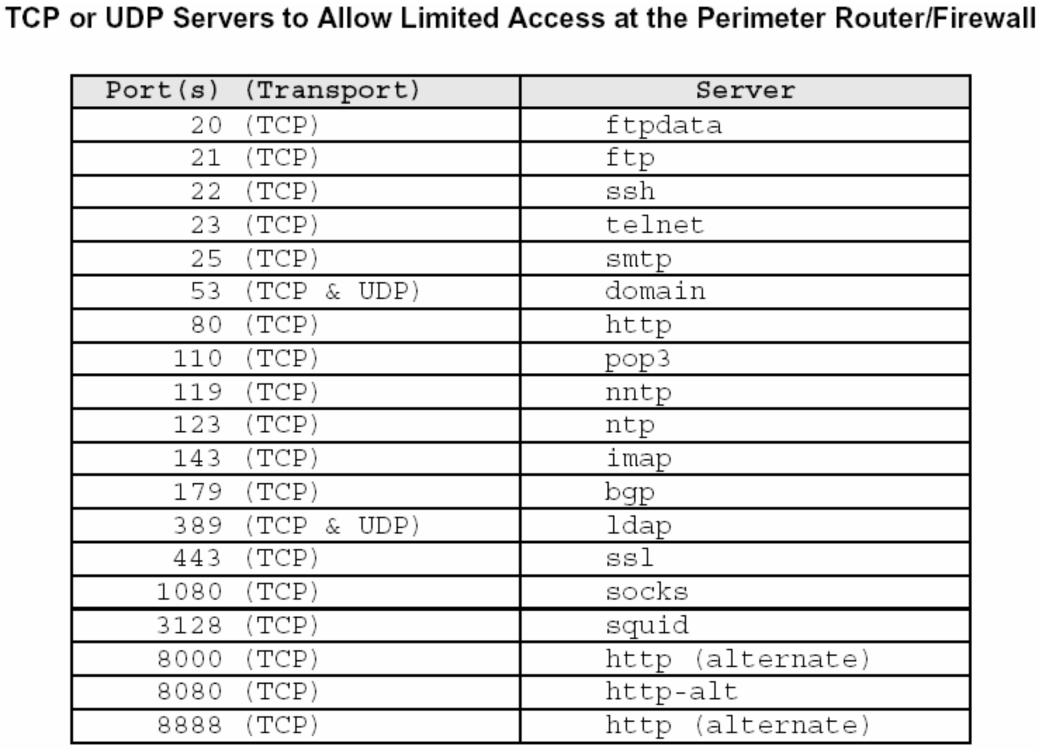

TCP or UDP Servers to Allow Limited Access at the Perimeter Router/Firewall

Port(s) (Transport)	Server
20 (TCP)	ftpdata
21 (TCP)	ftp
22 (TCP)	ssh
23 (TCP)	telnet
25 (TCP)	smtp
53 (TCP & UDP)	domain
80 (TCP)	http
110 (TCP)	pop3
119 (TCP)	nntp
123 (TCP)	ntp
143 (TCP)	imap
179 (TCP)	bgp
389 (TCP & UDP)	ldap
443 (TCP)	ssl
1080 (TCP)	socks
3128 (TCP)	squid
8000 (TCP)	http (alternate)
8080 (TCP)	http-alt
8888 (TCP)	http (alternate)

ICMP messages are commonly used by intruders as the basis for information gathering and denial of service attacks (see module 5, TCP/IP Security). It is highly recommended that special attention by given to restricting certain ICMP message types into and out of your network.

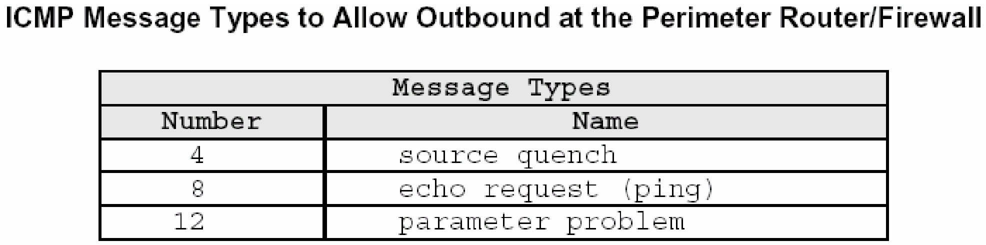

ICMP Message Types to Allow Outbound at the Perimeter Router/Firewall

Message Types	
Number	Name
4	source quench
8	echo request (ping)
12	parameter problem

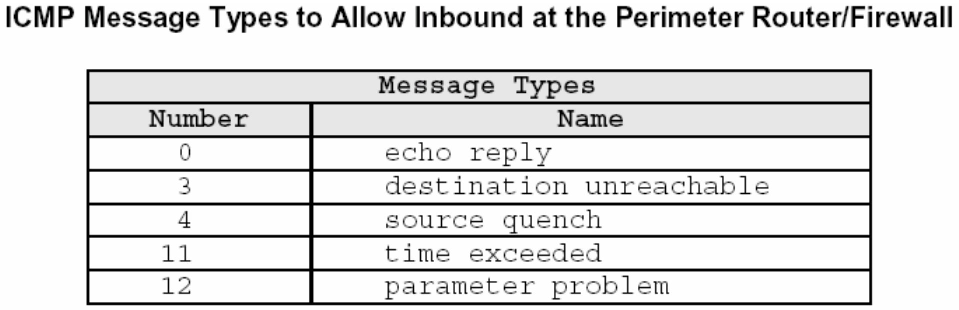

ICMP Message Types to Allow Inbound at the Perimeter Router/Firewall

Message Types	
Number	Name
0	echo reply
3	destination unreachable
4	source quench
11	time exceeded
12	parameter problem

An article on Cisco Press (http://www.cisco.com/public/cons/isp/essentials/Remote_Triggered_Black_Hole_Filtering-02.pdf) discusses Black Hole filtering.

Forwarding packets to Null 0 is a common way to filter packets to a specific destination. These are often done by creating specific static host routes and point them to the pseudo interface Null0. This technique commonly refereed as black hole routing or black hole filtering. Null0 is a pseudo-interface, which functions similarly to the null devices available on most operating systems. This interface is always up and can never forward or receive traffic. While Null0 is a pseudo interface, within CEF, it is not a valid interface. CEF (Cisco Express Forwarding) is an advanced, Layer 3 switching technology inside a router. It defines the fastest method by which a Cisco router forwards packets from ingress to egress interfaces. Hence, whenever a route is pointed to Null0, it will be dropped via CEF's and dCEF's.

Black Hole filtering uses the strength of the router's forwarding performance to drop black listed packets. A router's #1 job is forwarding packets - not filtering packets. The black hole routing technique uses the packet forwarding power to drop all packets bound for sites on the black list. In the ASIC forwarding world, this black holing has zero impact on the performance of the router (packets black holed to Null0 are cleared through a register clock). Software forwarding devices have some extra cycles needed to clear out a black-holed packet.

There are two main limitations with the black hole filtering technique. First, black hole filtering is Layer 3 only – not Layer 4. So access to all Layer 4 services at a give site will be blocked if black hole filtering is used. If selective Layer 4 filtering is necessary, use extended ACLs. For example, if you wish to drop all packets to a specific destination (i.e. an IP address), then black hole filtering is applicable. But, if you wish to drop all telnet packets to a destination, then black hole filtering is not applicable and an extended ACL is the optimum mitigation tool. Extended ACLs offer the fine Layer 4 granularity needed to filter at the application level. Second, it is hard to bypass or provide exceptions with the black hole filtering technique. Any organization that wishes to by-pass the black list must actually find a way to by-pass the filtering router's forward table. Compensation for either of these limitations are not trivial tasks. Yet, with due consideration and planning, options are available for both.

Remember, it's much easier to identify the specific traffic you want to allow, and then deny everything else. Following this practice will lessen the screening burden on your router, thereby freeing it up to do its primary function - routing data packets!

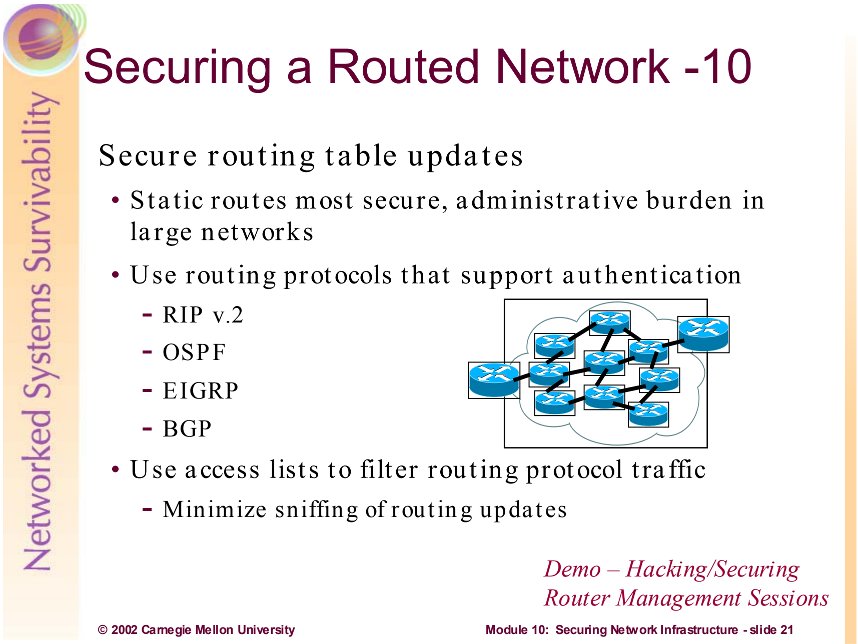

Networked Systems Survivability

Securing a Routed Network -10

Secure routing table updates

- Static routes most secure, administrative burden in large networks
- Use routing protocols that support authentication
  - RIP v.2
  - OSPF
  - EIGRP
  - BGP
- Use access lists to filter routing protocol traffic
  - Minimize sniffing of routing updates

Demo – Hacking/Securing Router Management Sessions

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 21

A router's primary responsibility is to send a packet of data to the intended destination. To accomplish this, each router needs a route table. Each router builds its table based on information from the network and from the network administrators. The router then uses a set of metrics, depending on the contents of the table and its routing algorithm, to compare routes and to determine the 'best' path to a destination.

Routers use four primary mechanisms for building their route tables:

- **Direct connection:** Any LAN segment to which the router is directly connected (hard wired from interface to interface) is automatically added to the route table.
- **Static routing:** A network administrator can manually instruct a router to use a given route to a particular destination. This method takes precedence over any other method of routing.
- **Dynamic routing:** Uses router update messages from other routers to create routes. The routing algorithm associated with the particular routing protocol determines the optimal path to a particular destination, and updates the route table. This method is the most flexible because it can automatically adapt to changes in the network.
- **Default routing:** Uses a manually entered route to a specific 'gateway of last resort' when route is not known by any other routing mechanism. This method is most useful for routers that serve as the sole connection between a small LAN and a large network like the Internet. Routers that depend on a single default gateway usually do not use routing protocols.

An unprotected router or routing domain makes an easy target for any network-savvy adversary. For example, an attacker who sends false routing update packets to an unprotected router can easily corrupt its route table. By doing this, the attacker can re-route network traffic in whichever manner he desires. The key to preventing such an attack is to protect the route tables from unauthorized and malicious changes. There are two basic approaches available for protecting route table integrity:

- **Use only static routes**: This may work in small networks, but is unsuitable for large networks due to the burden of administration.
- **Authenticate route table updates**: By using routing protocols with authentication, network administrators can deter attacks based on unauthorized routing changes. Authenticated router updates ensure that the update messages came from legitimate sources, bogus messages are automatically discarded.

Static routes are manually configured on the router as the only path to a given destination. These routes typically take precedence over routes chosen by dynamic routing protocols. In one sense, static routes are very secure. They are not vulnerable to spoofing attacks because they do not deal with router update packets. Exclusively using static routes will make network administration extremely difficult. Also, configuring a large network to use only static routes can make the availability of large pieces of the network subject to single points of failure. Static routes cannot handle events such as router failures. A dynamic routing protocol, however, such as OSPF, can correctly re-route traffic in the case of a router failure.

Because only the smallest networks can employ static routes, it is critical to select routing protocols that provide strong methods of neighbor authentication. The primary purpose of router neighbor authentication is to protect the integrity of a routing domain. In this case, authentication occurs when two neighboring routers exchange routing information. Authentication ensures that the receiving router incorporates into its tables only the route information that the trusted sending router really intended to send. It prevents a legitimate router from accepting and then employing unauthorized, malicious, or corrupted routing updates that would compromise the security or availability of a network. Such a compromise might lead to re-routing of traffic, a denial of service, or simply giving access to certain packets of data to an unauthorized person [NSA1].

All of the routing protocols listed on the slide support plaintext and cryptographic authentication. From a security standpoint, it is important to implement cryptographic authentication over simple plaintext authentication.

- **Plaintext authentication:** typically uses a shared plaintext key provided by the administrator. This key is inserted in the header of the routing protocol's packet. This is not a strong means of authentication, as this key can be sniffed by intruders, and then used as a basis of attack.
- **Cryptographic authentication:** typically utilizes a proven hashing algorithm (all of these use MD5) to encrypt a shared secret key provided by the administrator. Only the MD5 hash is transmitted between routers during routing updates.

Key Management is an important part of router authentication. It is recommended that one common key be used throughout your network (using multiple keys can be an administrative nightmare) and that distribution of that key be done manually. This key must be kept secret and should only be made known to a few administrators. The key should also be changed regularly.

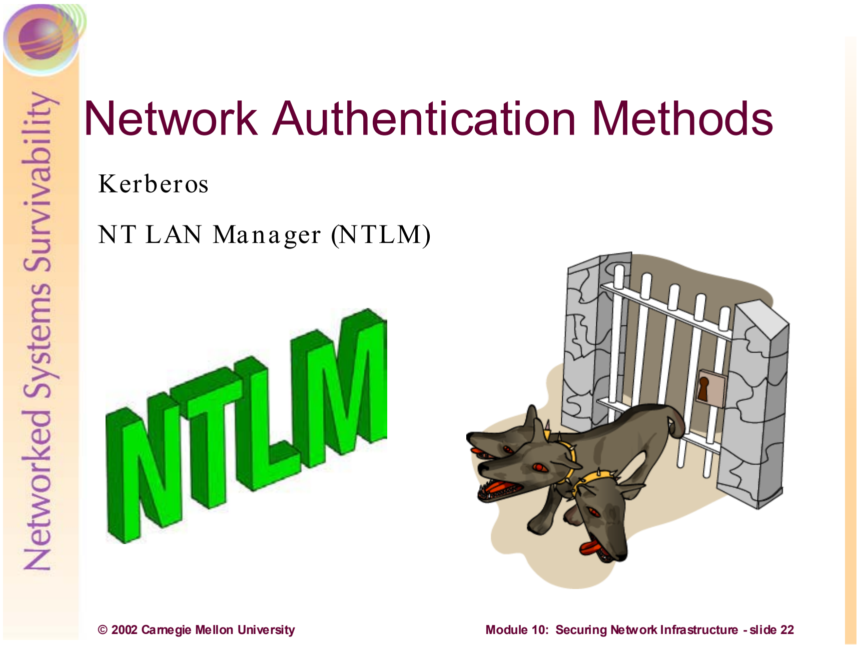

Networked Systems Survivability

Network Authentication Methods

- Kerberos
- NT LAN Manager (NTLM)

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 22

Modern computer systems provide service to multiple users and require the ability to accurately identify the user making a request. In traditional systems, the user's identity is verified by checking a password typed during login; the system records the identity and uses it to determine what operations may be performed. The process of verifying the user's identity is called authentication. Password based authentication is not suitable for use on computer networks. Passwords sent across the network can be intercepted and subsequently used by eavesdroppers to impersonate the user.

Authentication is the verification of the identity of a party who generated some data, and of the integrity of the data. A *principal* is the party whose identity is verified. The *verifier* is the party who demands assurance of the principal's identity. Data integrity is the assurance that the data received is the same as generated. Authentication mechanisms differ in the assurances they provide: some indicate that data was generated by the principal at some point in the past, a few indicate that the principal was present when the data was sent, and others indicate that the data received was freshly generated by the principal. Mechanisms also differ in the number of verifiers: some support a single verifier per message, while others support multiple verifiers. A third difference is whether the mechanism supports non-repudiation, the ability of the verifier to prove to a third party that the message originated with the principal.[27]

We will be covering the two most common means of network authentication in use today, Kerberos and NTLM. Kerberos is widely used on heterogeneous networks with multiple operating systems. NTLM is used on Microsoft Windows based networks.

[27] http://www.isi.edu/gost/publications/kerberos-neuman-tso.html

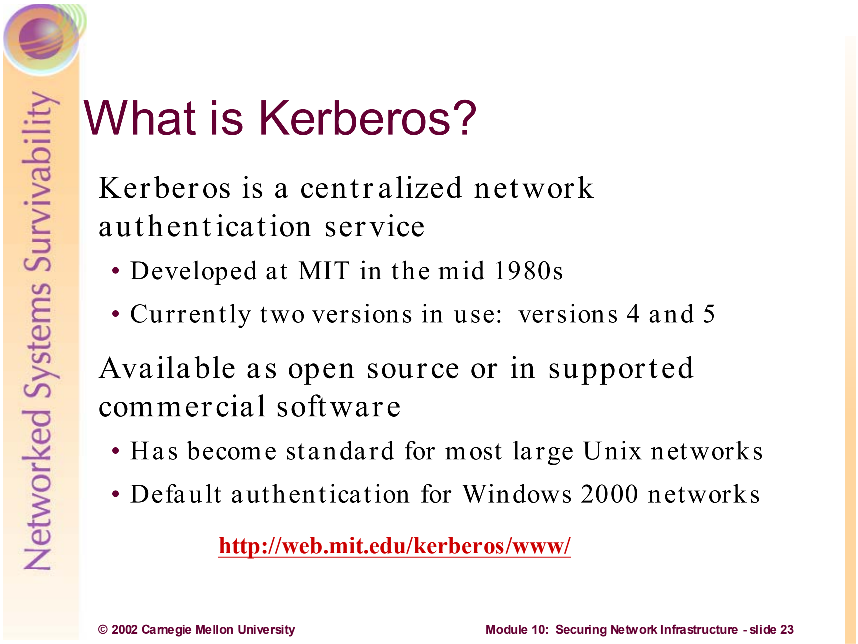

Networked Systems Survivability

What is Kerberos?

Kerberos is a centralized network authentication service

- Developed at MIT in the mid 1980s
- Currently two versions in use: versions 4 and 5

Available as open source or in supported commercial software

- Has become standard for most large Unix networks
- Default authentication for Windows 2000 networks

http://web.mit.edu/kerberos/www/

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 23

Kerberos is an authentication service developed at the Massachusetts Institute of Technology (MIT). It's purpose is to allow users and services to *authenticate* themselves to each other. That is, it allows them to demonstrate their identity to each other, unequivocally (it is hoped).[28] The name Kerberos stems from the mythological 3-headed dog that guards the gates of Hades.

Kerberos was designed by MIT in the 1980s as part of Project Athena. Project Athena was an initiative to investigate how computer technology could integrate into an undergraduate curriculum. It lasted from 1983-1991. Many computing technologies were researched and developed, resulting in MIT’s current distributed computing environment (DCE). Kerberos was the authentication technology designed for MIT’s DCE.

There are currently two versions in use: versions 4 and 5. While still supported, MIT considers version 4 to be dead and is concentrating all development efforts on version 5. Therefore, version 5 has become the defacto standard for Kerberos. See RFC 1510[29] for the complete specification of Kerberos version 5.

MIT freely distributes source code for Kerberos. It is also widely used by vendors of networking software—most notably, Microsoft has incorporated Kerberos version 5 into Windows 2000. It is the default authentication methodology on all Windows 2000 networks.[30]

[28] http://www.isi.edu/gost/brian/security/kerberos.html

[29] ftp://ftp.isi.edu/in-notes/rfc1510.txt

[30] http://www.microsoft.com/windows2000/docs/kerberos.doc

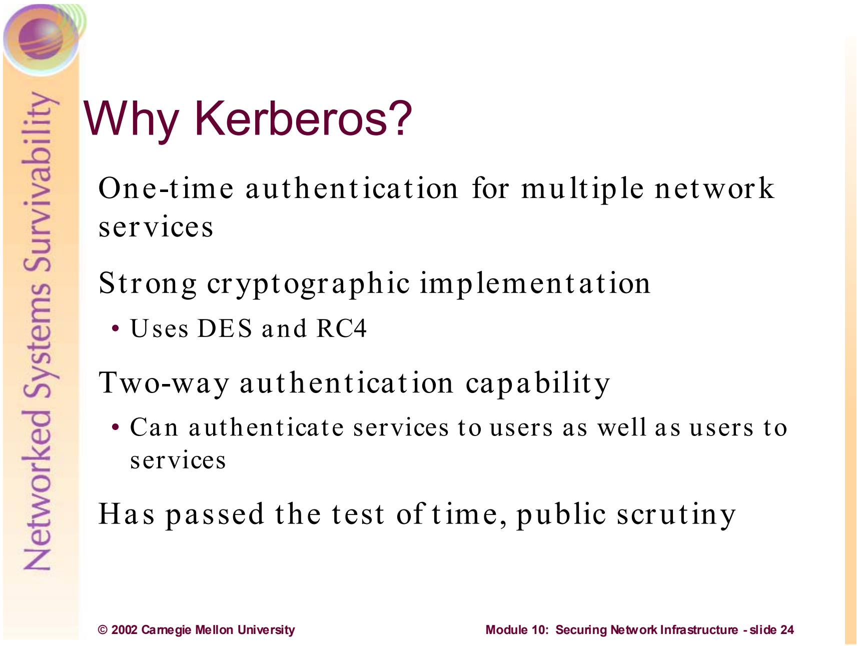
Networked Systems Survivability

Why Kerberos?

- One-time authentication for multiple network services
- Strong cryptographic implementation
  - Uses DES and RC4
- Two-way authentication capability
  - Can authenticate services to users as well as users to services
- Has passed the test of time, public scrutiny

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 24

Kerberos provides for single sign-on for accessing multiple network services. This means users need not authenticate themselves every single time they access a different network service.

Fundamental to Kerberos is that no user passwords are ever passed in plaintext over the network. To do this, Kerberos utilizes public cryptographic algorithms. In the MIT implementation, DES and Triple DES are used for encryption and MD4 or SHA-1 are used for its Hash Message Authentication Code (HMAC). Microsoft's implementation uses the RC4 algorithm for encryption and MD5 as its HMAC.[31] This presents some interoperability issues when working across both implementations,[32] however, Kerberos provides enough flexibility for almost all heterogeneous environments.

One of the strengths of Kerberos is that it provides for 2-way authentication. That is, servers can verify users and vice versa.

Kerberos is an open source standard, meaning its source code is freely available for anyone to inspect. The wide distribution of Kerberos has contributed to its success, with features and security refinements being included as a result of various real-world implementations. It has passed the test of public scrutiny and is generally accepted in the security community to be "as good as it gets" when it comes to network authentication.

[31] http://www.ietf.org/internet-drafts/draft-brezak-win2k-krb-rc4-hmac-03.txt

[32] http://www.microsoft.com/windows2000/techinfo/planning/security/kerbsteps.asp

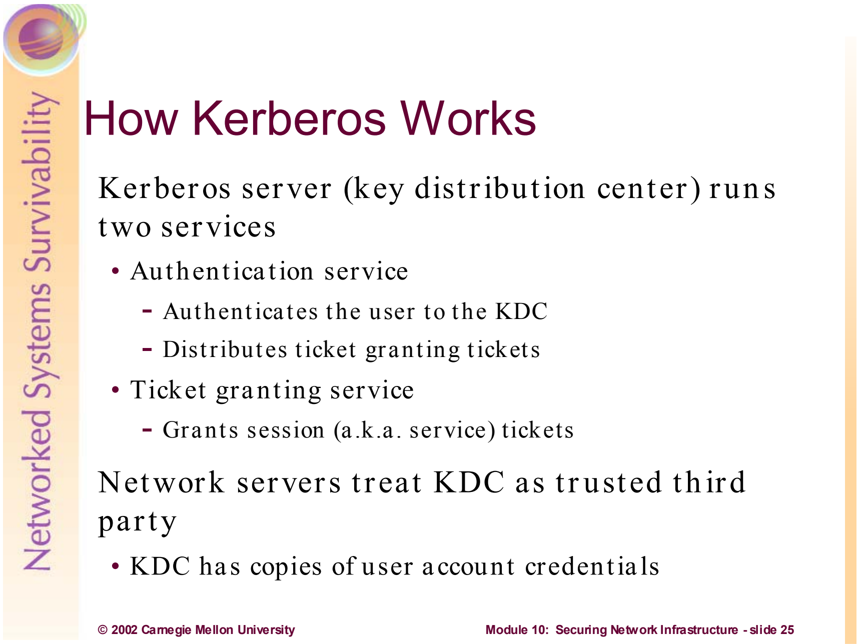

Networked Systems Survivability

How Kerberos Works

- Kerberos server (key distribution center) runs two services
  - Authentication service
    - Authenticates the user to the KDC
    - Distributes ticket granting tickets
  - Ticket granting service
    - Grants session (a.k.a. service) tickets
- Network servers treat KDC as trusted third party
  - KDC has copies of user account credentials

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 25

The Kerberos Server is called a key distribution center (KDC). This is the centralized authentication server for Kerberos. The KDC operates two services:

- **Authentication Service (AS)**: checks the authenticity of the user's inputted credentials at initial logon. If these credentials are validated against the KDC's database, the authentication service returns a Ticket Granting Ticket (TGT) to the user. The TGT is then sent to the ticket granting service when requesting access to a network resource.
- **Ticket Granting Service (TGS)**: Supplies sessions tickets (ST) to clients after receiving TGT

Much the same way Verisign Inc. is trusted to verify the validity of an e-commerce web site's certificate for an Internet shopper, the KDC is treated similarly by network clients and servers—as a trusted third party.

As part of its requirement to authenticate users and servers, the KDC must maintain a database of user credentials and there keys which are cryptographically derived from their passwords. The database must also contain the keys of all kerberized network servers.

Because the Kerberos authentication process is rather complicated, a series of illustrations will be used to explain how it works. A very good introductory (this one is actually fun to read!) tutorial on Kerberos is available at: http://web.mit.edu/kerberos/www/dialogue.html

Following are simplified illustrations of how the 2 Kerberos services work—using Windows 2000 as an example:

Kerberos Authentication is used when a Windows 2000 client is authenticated by a Windows 2000 domain controller. A Windows 2000 domain controller functions as a Kerberos Key Distribution Center or KDC. The KDC runs two services, an Authentication Service and a Ticket Granting Service. When a user provides credentials to the KDC, the Authentication Service verifies that the user is legitimate and sends the user a ticket granting ticket. A ticket granting ticket is like the pass you get when you pay your way into an amusement park. Once inside, you need to show the pass in order to ride the rides. With Kerberos, you actually show the pass to the Ticket Granting Service to request a ticket for a resource.

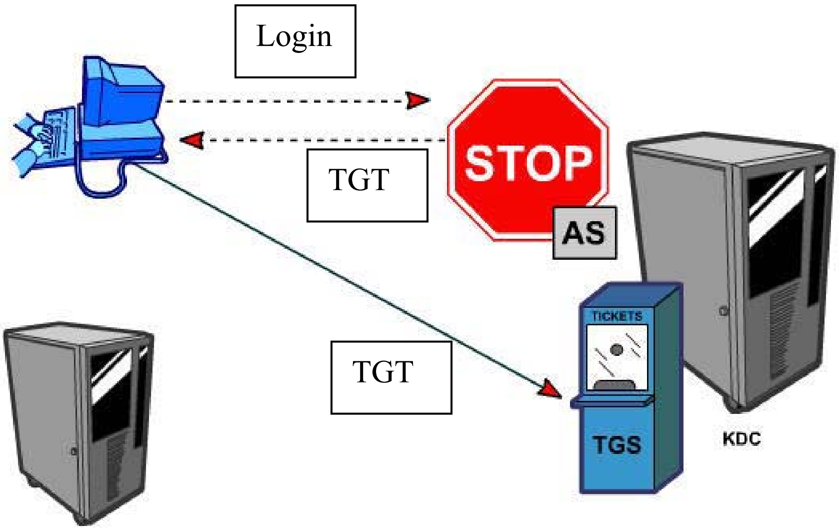

The TGS grants you a ticket that can be used for a predetermined amount of time. The key is encrypted with a key known only to the TGS and the server that stores the requested resource. The user will present this ticket to the server where the resource is located. The server will decrypt the ticket and look at the time stamp to make sure the ticket was issued during a valid time period. This is similar to the color coding some amusement parks use to make sure that visitors don't use tickets that were thrown out the day before.

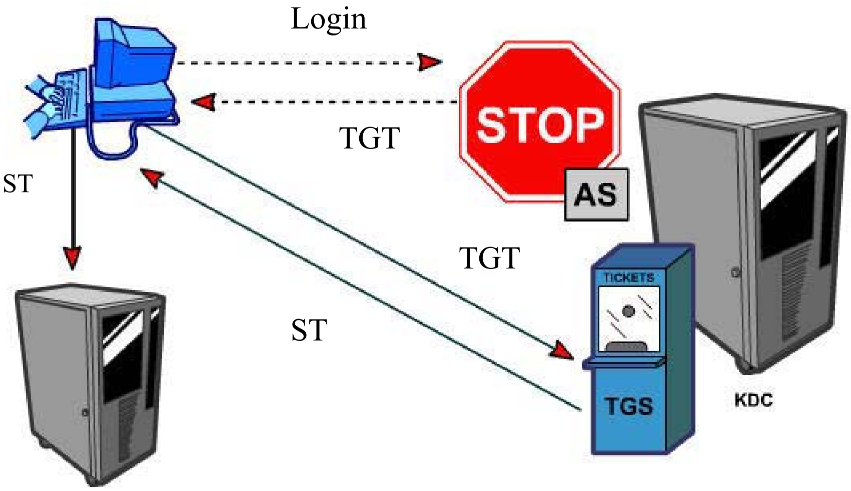

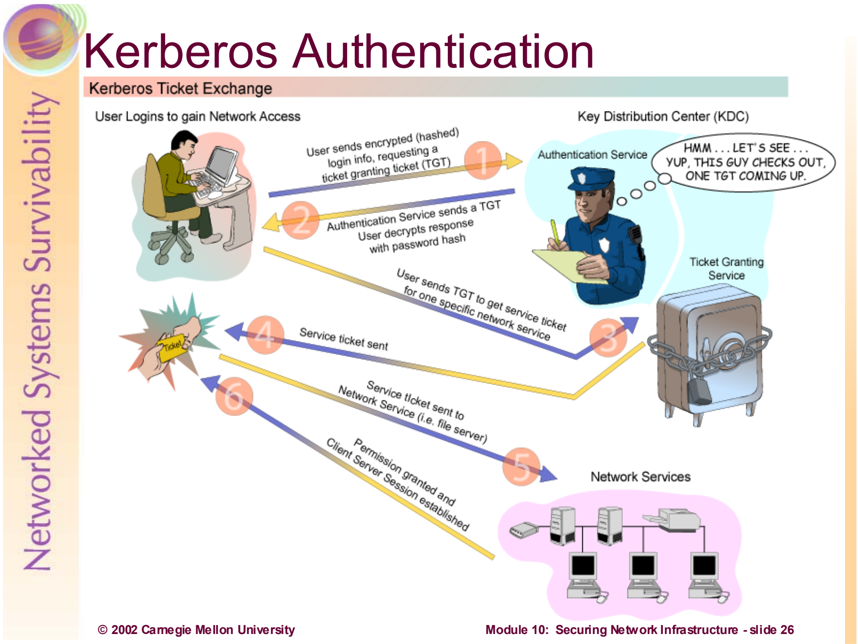

Kerberos Authentication

Kerberos Ticket Exchange

Networked Systems Survivability
User Logins to gain Network Access
Key Distribution Center (KDC)
User sends encrypted (hashed) login info, requesting a ticket granting ticket (TGT)
1
Authentication Service
HMM . . . LET'S SEE . . . YUP, THIS GUY CHECKS OUT, ONE TGT COMING UP.
2
Authentication Service sends a TGT User decrypts response with password hash
Ticket Granting Service
User sends TGT to get service ticket for one specific network service
3
Ticket
4
Service ticket sent
Service ticket sent to Network Service (i.e. file server)
5
6
Permission granted and Client Server Session established
Network Services

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 26

Following are step by step (more detailed) illustrations of the Kerberos authentication process:

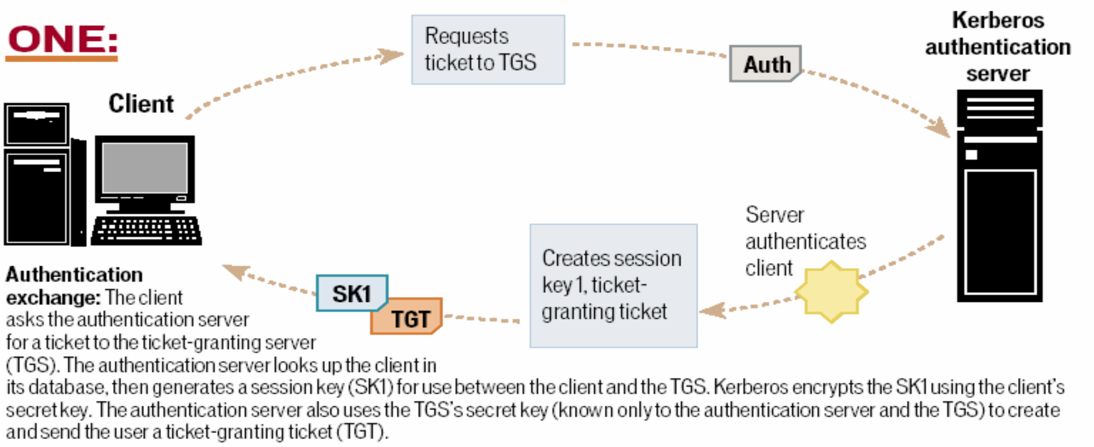

ONE:
Client
Requests ticket to TGS
Auth
Kerberos authentication server
Server authenticates client
Creates session key 1, ticket-granting ticket
SK1
TGT

Authentication exchange: The client asks the authentication server for a ticket to the ticket-granting server (TGS). The authentication server looks up the client in its database, then generates a session key (SK1) for use between the client and the TGS. Kerberos encrypts the SK1 using the client's secret key. The authentication server also uses the TGS's secret key (known only to the authentication server and the TGS) to create and send the user a ticket-granting ticket (TGT).

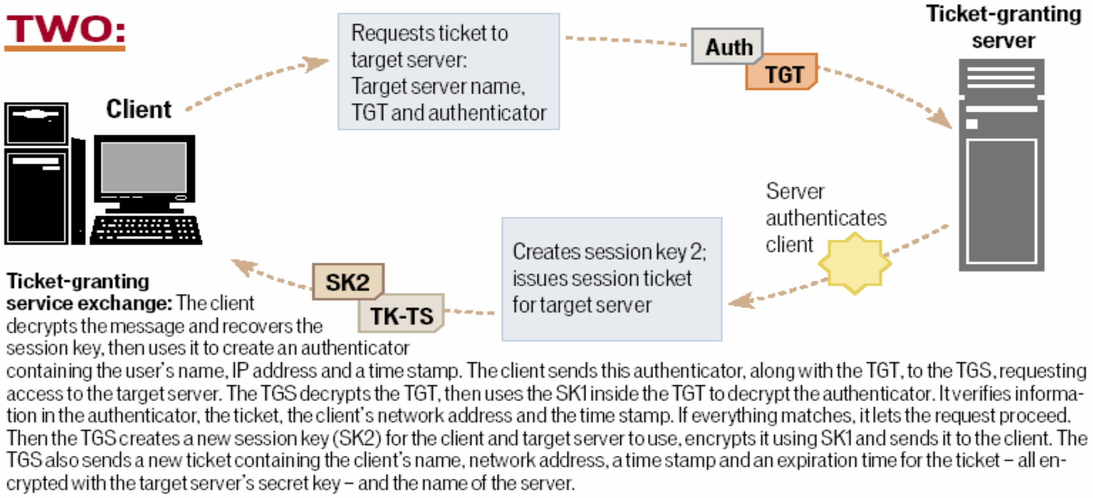

TWO:
Client
Requests ticket to target server: Target server name, TGT and authenticator
Auth
TGT
Ticket-granting server
Server authenticates client
Creates session key 2; issues session ticket for target server
SK2
TK-TS

Ticket-granting service exchange: The client decrypts the message and recovers the session key, then uses it to create an authenticator containing the user's name, IP address and a time stamp. The client sends this authenticator, along with the TGT, to the TGS, requesting access to the target server. The TGS decrypts the TGT, then uses the SK1 inside the TGT to decrypt the authenticator. It verifies information in the authenticator, the ticket, the client's network address and the time stamp. If everything matches, it lets the request proceed. Then the TGS creates a new session key (SK2) for the client and target server to use, encrypts it using SK1 and sends it to the client. The TGS also sends a new ticket containing the client's name, network address, a time stamp and an expiration time for the ticket – all encrypted with the target server's secret key – and the name of the server.

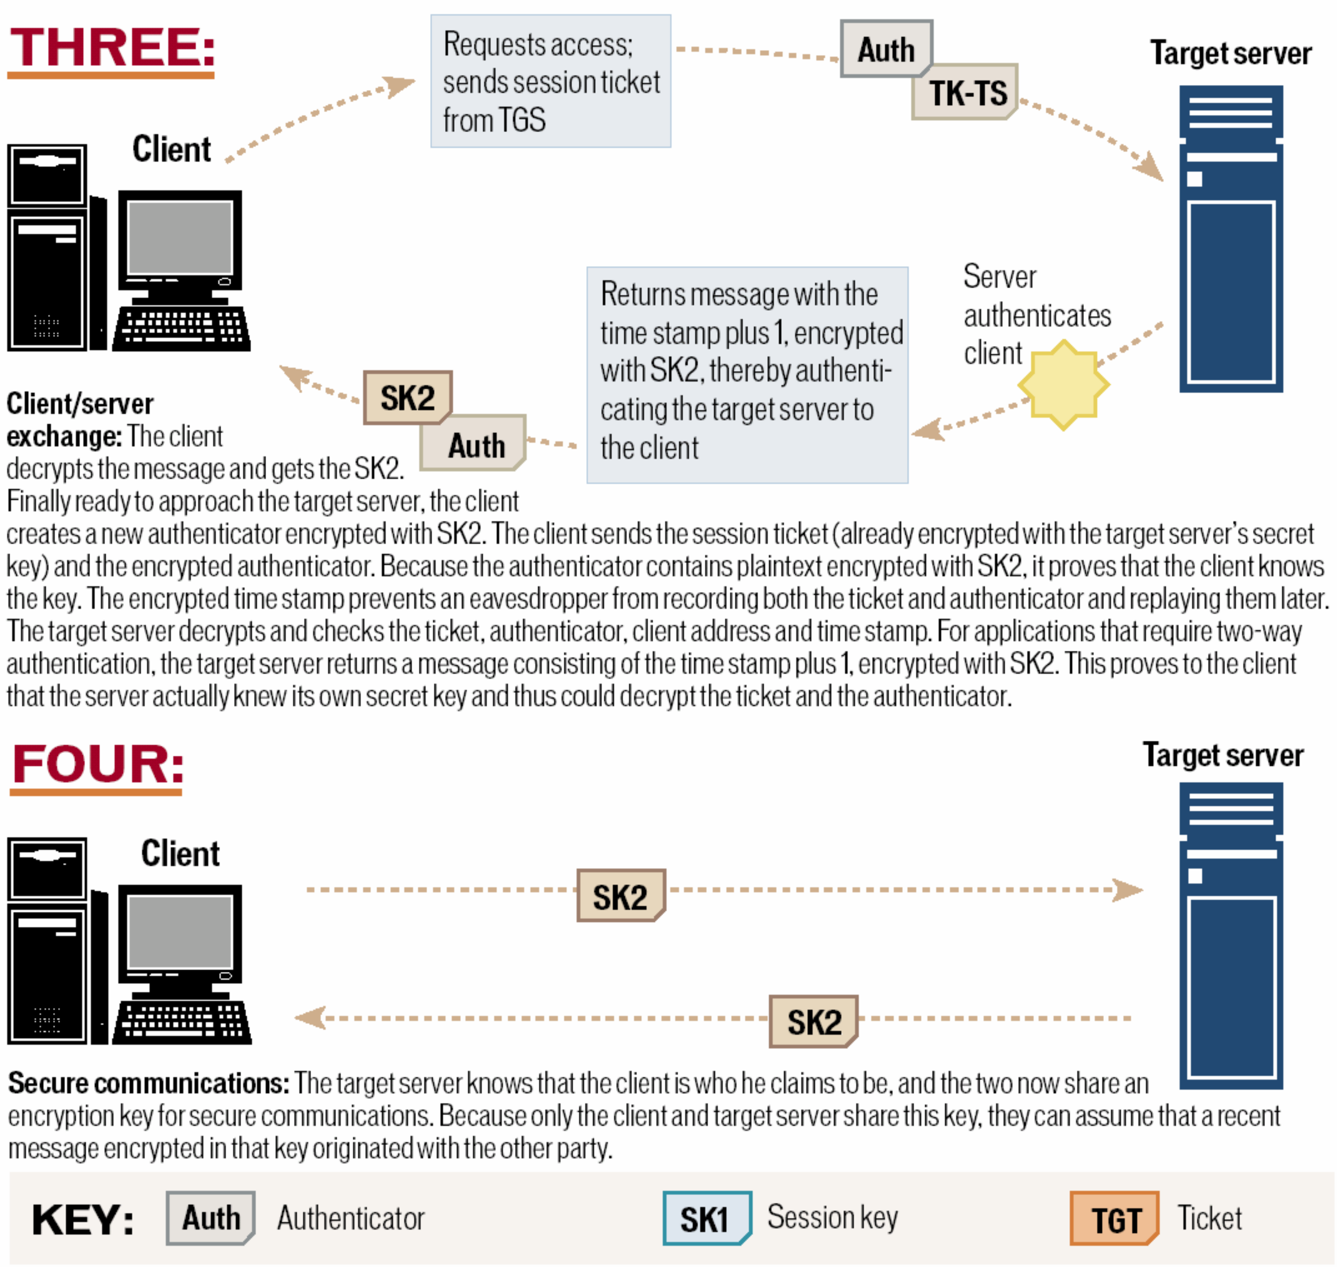
THREE:
Client
Requests access; sends session ticket from TGS
Auth
TK-TS
Target server
Returns message with the time stamp plus 1, encrypted with SK2, thereby authenticating the target server to the client
Server authenticates client
SK2
Auth
FOUR:
Client
Target server
SK2
SK2
KEY:
Auth
Authenticator
SK1
Session key
TGT
Ticket

Client/server exchange: The client decrypts the message and gets the SK2. Finally ready to approach the target server, the client creates a new authenticator encrypted with SK2. The client sends the session ticket (already encrypted with the target server's secret key) and the encrypted authenticator. Because the authenticator contains plaintext encrypted with SK2, it proves that the client knows the key. The encrypted time stamp prevents an eavesdropper from recording both the ticket and authenticator and replaying them later. The target server decrypts and checks the ticket, authenticator, client address and time stamp. For applications that require two-way authentication, the target server returns a message consisting of the time stamp plus 1, encrypted with SK2. This proves to the client that the server actually knew its own secret key and thus could decrypt the ticket and the authenticator.

Secure communications: The target server knows that the client is who he claims to be, and the two now share an encryption key for secure communications. Because only the client and target server share this key, they can assume that a recent message encrypted in that key originated with the other party.

Source: http://www.computerworld.com/computerworld/records/images/pdf/kerberos_chart.pdf

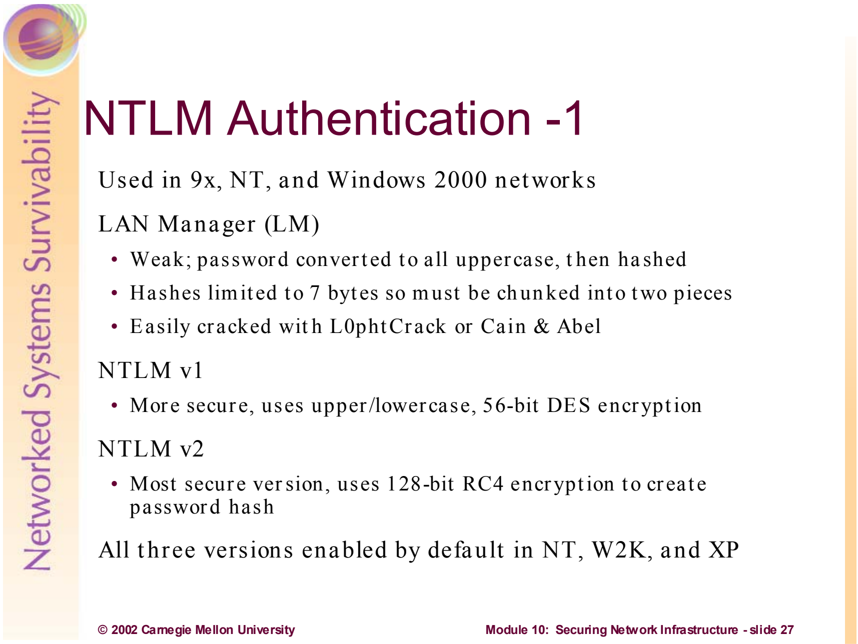

Networked Systems Survivability

NTLM Authentication -1

Used in 9x, NT, and Windows 2000 networks

LAN Manager (LM)

- Weak; password converted to all uppercase, then hashed
- Hashes limited to 7 bytes so must be chunked into two pieces
- Easily cracked with L0phtCrack or Cain & Abel

NTLM v1

- More secure, uses upper/lowercase, 56-bit DES encryption

NTLM v2

- Most secure version, uses 128-bit RC4 encryption to create password hash

All three versions enabled by default in NT, W2K, and XP

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 27

NTLM is a challenge response authentication method utilized exclusively on Microsoft Windows networks.

NTLM has three different versions—all enabled by default on NT/2000 networks:

- **Lan Manager (LM)**: is utilized for authenticated legacy Windows 3.1, 9x systems
  - This should be disabled on all Windows networks as it is easily cracked
  - If Windows 9x systems reside on your network, Directory Client Services (available on Windows 2000 CD) must be installed on these systems in order to allow NTLM v2 authentications.
  - To disable LanMan authentication, set the following registry key:

    Hive: HKEY_LOCAL_MACHINE

    Key: System\CurrentControlSet\Control\Lsa

    Name: LMCompatibilityLevel

    Type: REG_DWORD

    Value: 5
- NTLM v1 is better than LM, however it still is susceptible to attack
- NTLM v2 is preferred version—use this version and disable the other 2
  - In Windows 2000, can be done by utilizing the Local Security Policy snap in:

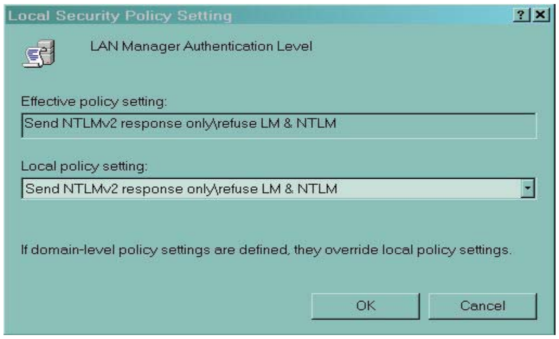

Local Security Policy Setting

LAN Manager Authentication Level

Effective policy setting:

Send NTLMv2 response only\refuse LM & NTLM

Local policy setting:

Send NTLMv2 response only\refuse LM & NTLM

If domain-level policy settings are defined, they override local policy settings.

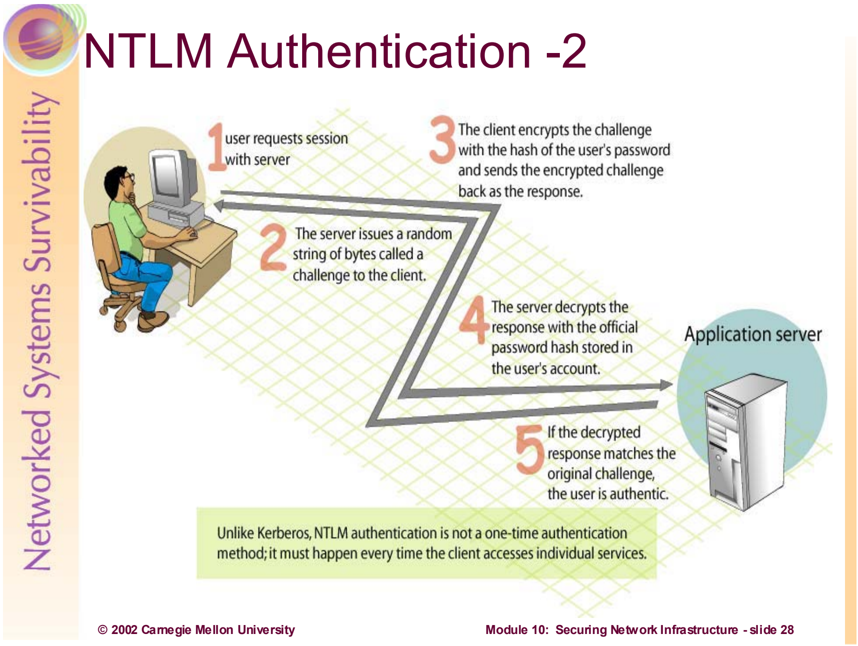

NTLM Authentication -2

Networked Systems Survivability

1 user requests session with server

2 The server issues a random string of bytes called a challenge to the client.

3 The client encrypts the challenge with the hash of the user's password and sends the encrypted challenge back as the response.

4 The server decrypts the response with the official password hash stored in the user's account.

5 If the decrypted response matches the original challenge, the user is authentic.

Unlike Kerberos, NTLM authentication is not a one-time authentication method; it must happen every time the client accesses individual services.

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 28

Here is how NTLM Challenge/Response works:

When a user needs to connect to a server, the server authenticates the user with a challenge/response protocol. The server issues a random string of bytes called a *challenge* to the client. The client encrypts the challenge with the hash of the user's password and sends the encrypted challenge back as the response. The server decrypts the response with the official password hash stored in the user's account. If the decrypted response matches the original challenge, the user is authentic.

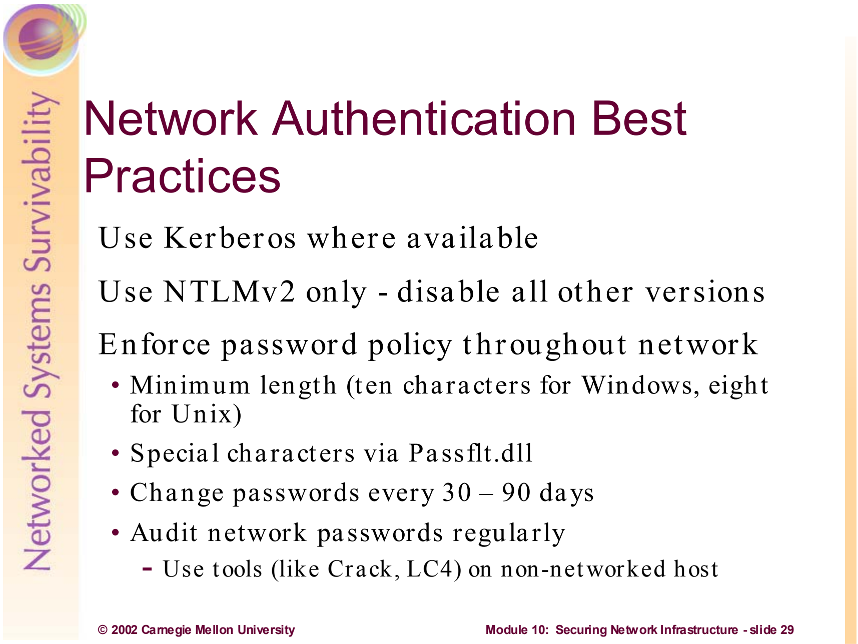

Networked Systems Survivability

Network Authentication Best Practices

Use Kerberos where available

Use NTLMv2 only - disable all other versions

Enforce password policy throughout network

- Minimum length (ten characters for Windows, eight for Unix)
- Special characters via Passflt.dll
- Change passwords every 30 – 90 days
- Audit network passwords regularly
  - Use tools (like Crack, LC4) on non-networked host

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 29

As Kerberos is the strongest network authentication method available today, it should be utilized exclusively where available. The key word here is exclusively—the security of your networks will suffer if you allow multiple authentication methods on your network. If you are on a Windows network and must allow NTLM for non-kerberized (9x and NT) clients, use NTLM v2 only.

A summary Windows example of Kerberos and NTLM authentication is illustrated below:

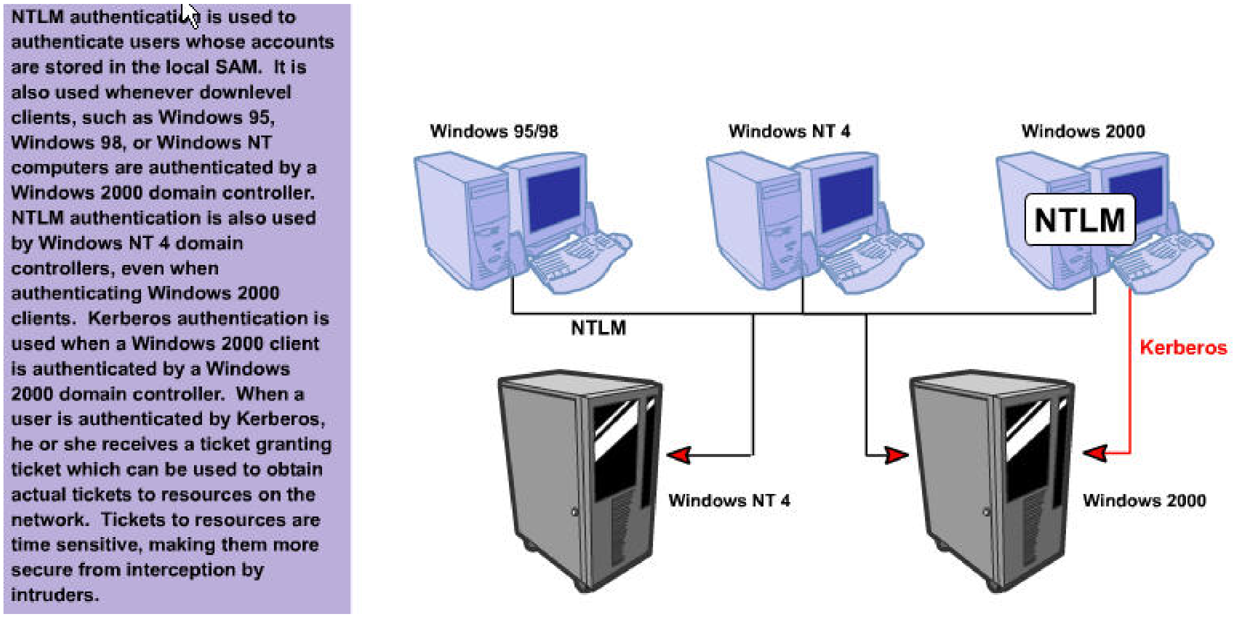

NTLM authentication is used to authenticate users whose accounts are stored in the local SAM. It is also used whenever downlevel clients, such as Windows 95, Windows 98, or Windows NT computers are authenticated by a Windows 2000 domain controller. NTLM authentication is also used by Windows NT 4 domain controllers, even when authenticating Windows 2000 clients. Kerberos authentication is used when a Windows 2000 client is authenticated by a Windows 2000 domain controller. When a user is authenticated by Kerberos, he or she receives a ticket granting ticket which can be used to obtain actual tickets to resources on the network. Tickets to resources are time sensitive, making them more secure from interception by intruders.

Poor password selection is frequently a major problem for any system's security. Users should be forced to change their passwords regularly. Set up password aging via Account Policy for Windows systems or the /etc/default/passwd file in Unix. Administrators should obtain and run password-guessing programs (i.e., “John the Ripper,'' “L0phtCrack,” and “Crack”) frequently to identify those users having easily guessed passwords. Because password cracking programs are very CPU intensive and can slow down the system on which it is running, it is a good idea to transfer the encrypted passwords (the dumped SAM database for Windows and the /etc/passwd and /etc/shadow files in Unix) to a stand-alone (not networked) system. Also, by doing the work on a non-networked machine, any results found will not be accessible by anyone unless they have physical access to that system. As a general rule, passwords should:

- Be 12 or more characters in length on Windows systems, 8 characters in length on Unix
- Include upper and lower case letters, numbers, and special characters—enforced in Windows by implementing passflt.dll
- Not consist of dictionary words
- Be changed regularly (every 30 to 90 days)
- For Unix, be encrypted and stored in the /etc/shadow file (for some Unix systems) with permissions set to 400 with ownership by root and group sys. The /etc/passwd file should have permissions 644 with owner root and group root.
- Be cracked every month to find users choosing easily guessed or cracked passwords

For Unix, lock the following accounts by placing a *LK* in encrypted password field in /etc/shadow: adm, bin, daemon, listen, lp, nobody, noaccess, nuucp, smtp, sys, uucp. These accounts should not have login shells, rather they should be set to /dev/null [NSA2].

*****It's important to remember that even Kerberos and NTLMv2 can be undermined by weak passwords.*****

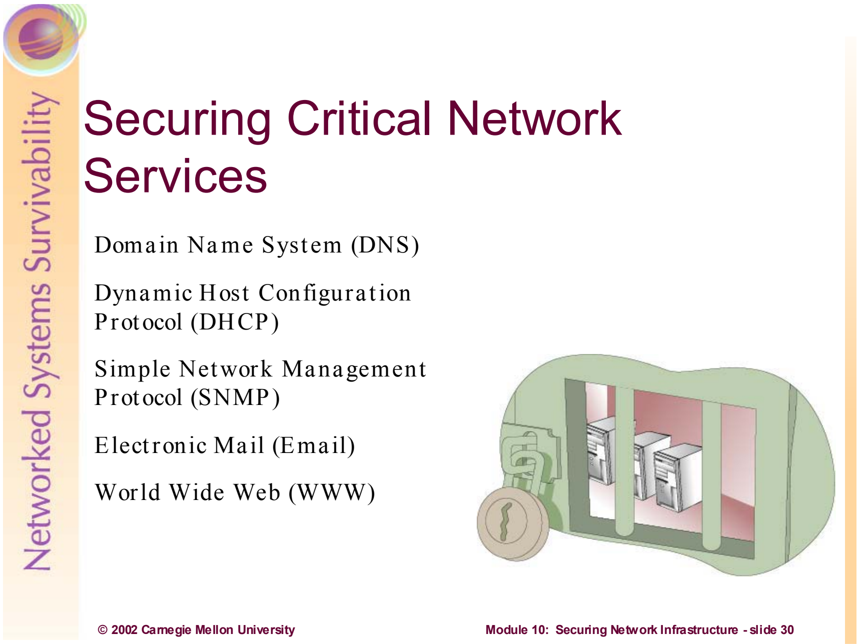
Networked Systems Survivability

Securing Critical Network Services

- Domain Name System (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- Simple Network Management Protocol (SNMP)
- Electronic Mail (Email)
- World Wide Web (WWW)

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 30

Key network services have become vital to the missions of most organizations today—so much so that prolonged interruptions can have a drastic (if not catastrophic) effect on the survivability of the organization.

Many of these services were originally designed in the early days of the Internet, when connectivity was the concern—not security. This fact makes securing these services rather challenging for administrators. We will cover some security best practices that can help administrators with this challenge.

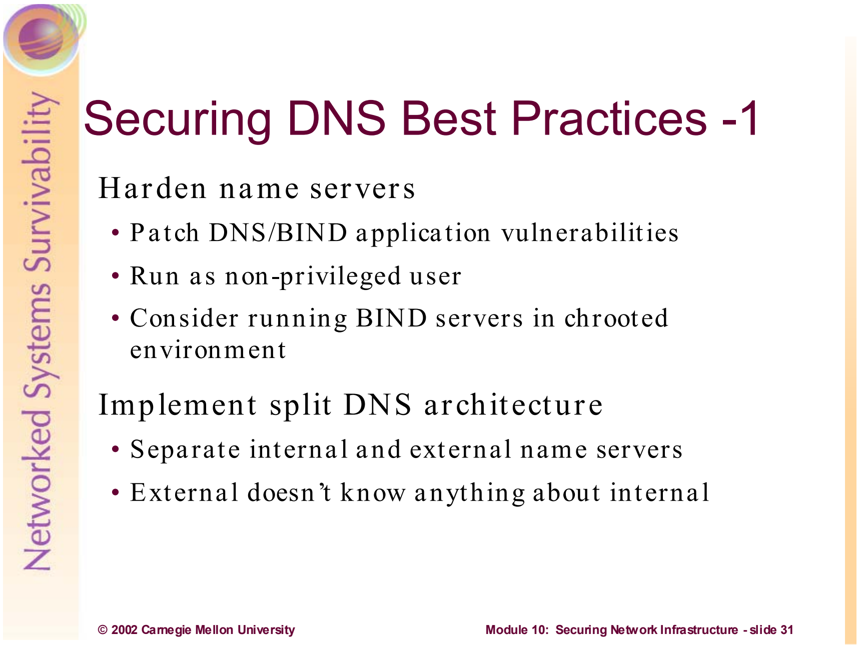

Networked Systems Survivability

Securing DNS Best Practices -1

Harden name servers

- Patch DNS/BIND application vulnerabilities
- Run as non-privileged user
- Consider running BIND servers in chrooted environment

Implement split DNS architecture

- Separate internal and external name servers
- External doesn't know anything about internal

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 31

The Domain Name System is the global, hierarchical, distributed database that resolves host names (like www.cert.org) to IP addresses (192.88.209.14). A working knowledge of how DNS works is assumed for this course, however a good summary of this (for review purposes) is available at: http://www.garykessler.net/library/dns.html

Because name servers are so critical to the operations of modern computer networks, it is important to make them a secure as possible. Best practices for doing this are:

- Harden the underlying Unix[33] or Microsoft[34] operating system of the server[35]
- Review the version of DNS or Berkeley Internet Name Domain (BIND) server application you are using for known vulnerabilities and then apply appropriate patches or workarounds. CERT advisories and vulnerability notes[36] should be consulted.
- If at all possible, run the DNS/BIND server application with as few privileges as possible. That way if the server is compromised, it will limit the damage that can be done to other network resources.
- Consider running BIND application in a chroot environment. Chroot is a Unix command used to run a command or interactive shell with a special root directory. It can be used to create a "virtual" operating system and directory tree and can enhance security.[37]
- On Unix servers, disable the BIND name daemon (named) on all Unix systems not authorized to be name servers—on required BIND servers, hide the version string via the version option in named.conf. This will make it more difficult for intruders to enumerate what kind of server you are running [NSA2].

A fundamental best practice for securing DNS is to implement a split DNS architecture. This means that two independent DNS servers are implemented, one outside of the protected network (in a DMZ) and one on the inside. The internal server contains the database of all the DNS names within the organization,

[33] http://www.cert.org/tech_tips/usc20_full.html
[34] http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/tools/tools.asp
[35] http://www.cert.org/security-improvement/modules/m10.html
[36] http://www.cert.org/nav/index_red.html
[37] http://www.linuxdoc.org/HOWTO/Chroot-BIND-HOWTO.html

whereas the external server knows only how to resolve names dealing with the external presence, such as e-mail forwarders and web servers.

The External DNS Server resides in the DMZ and is for public use. The only information that should be on the External DNS Server is the information that needs to be advertised to Internet clients. No internal information should be available to the External DNS Server. Furthermore, only the External DNS Server can communicate with the outside. The Internal DNS Server should handle resolution of external DNS information for the firewall and all hosts on the Protected Network. Although the Internal DNS Server is unable to communicate directly with the external network, it should be configured to send queries and receive the responses via the External DNS Server. For proper functionality, DNS UDP connections should be allowed through the firewall to the External DNS Server. These connections allow external clients to query the DNS. However, DNS TCP connections (needed for zone transfers) should only be allowed between the External DNS server and trusted DNS servers located beyond the External Router. All other DNS TCP connections should be denied. This setting reduces an attacker's ability to map the Protected Network via zone transfers. Also, note that even if DNS information is hidden there are other sources that will provide information about the internal naming scheme. For instance, email headers, NetBIOS, and other services could still supply host information to an attacker.

The Internal DNS Server provides the functionality required for registering services and clients in the domain (Windows) and for internal name resolution. This server is extremely important because it holds a complete map of the domain. Therefore, this server should not share any information with the External DNS Server. Since there may be requirements for reverse lookup zones to include clients in the Protected Network, the reverse lookup zone on the internal server may need to be passed to the External DNS Server [NSA3].

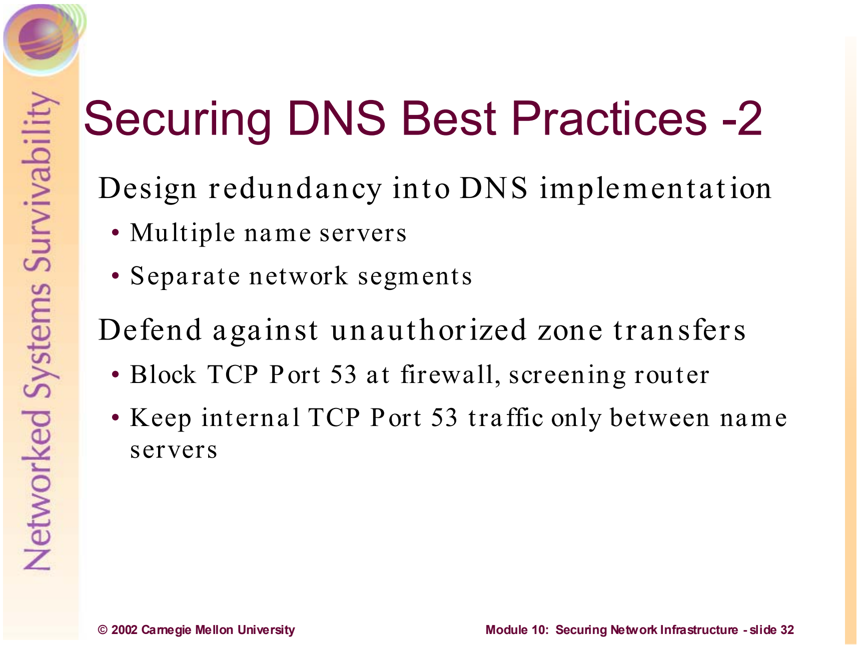

Networked Systems Survivability

Securing DNS Best Practices -2

Design redundancy into DNS implementation

- Multiple name servers
- Separate network segments

Defend against unauthorized zone transfers

- Block TCP Port 53 at firewall, screening router
- Keep internal TCP Port 53 traffic only between name servers

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 32

Redundancy and fault tolerance for an organization's DNS is crucial to survivability:

- Implement at least 2 internal and external name servers; and consider implementing some form of fault tolerant clustering for these servers
- It is advisable to place all of these servers on separate physical network segments, thereby mitigating potential single points of failure based on network infrastructure problems.

Zone transfers ensure that all DNS servers have the same information in their zone file. Maliciously attempted zone transfers are one of the most common information gathering attacks found on the Internet today. Some DNS server software includes methods for authenticating zone transfers between authorized name servers. These can be proprietary methods like those internal to Windows 2000 DNS,[38] or they can be implementations of TSIG and or DNSec (covered next slide). It is highly recommended that zone transfer authentication be utilized wherever possible.

DNS traffic travels through port 53 (UDP and TCP). UDP port 53 is required for client queries while the TCP port is required for zone transfers. In most cases, it is unnecessary to allow zone transfers outside of the Protected Network so TCP port 53 should be blocked at the firewall, and screening routers. Internally, access lists on routers should be configured so that TCP port 53 traffic can flow only between name servers.

[38] http://nsa1.www.conxion.com/win2k/guides/w2k-6.pdf

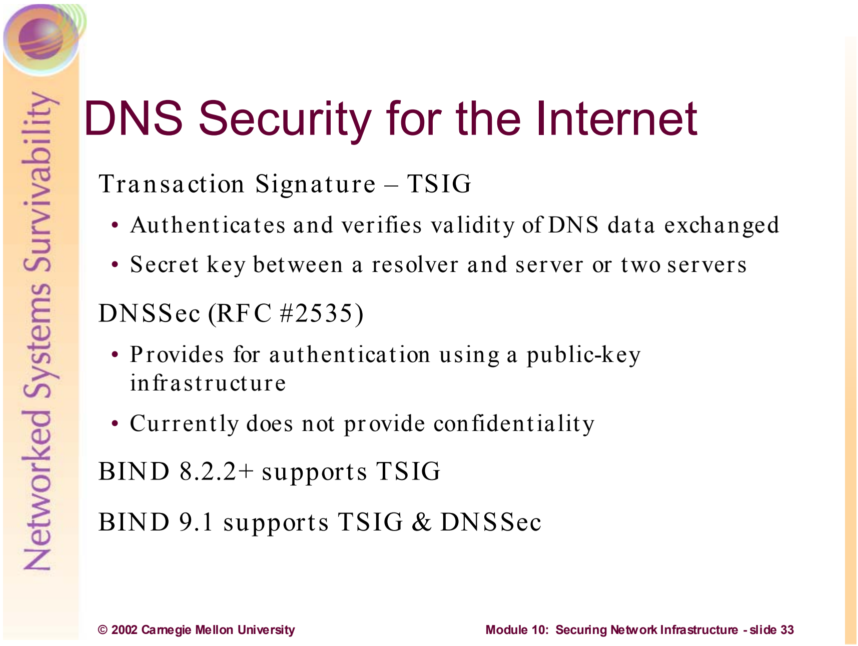
Networked Systems Survivability

DNS Security for the Internet

Transaction Signature – TSIG

- Authenticates and verifies validity of DNS data exchanged
- Secret key between a resolver and server or two servers

DNSSec (RFC #2535)

- Provides for authentication using a public-key infrastructure
- Currently does not provide confidentiality

BIND 8.2.2+ supports TSIG

BIND 9.1 supports TSIG & DNSSec

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 33

As mentioned on the previous page, authenticating DNS traffic is highly desirable from a security standpoint. There are currently 2 IETF sponsored techniques for doing this, TSIG and DNS Security Extensions. It is important to note that neither of these specifications incorporate any form of data encryption (like IPSec); they only perform authentication. Therefore, plaintext DNS data could still be sniffed on the wire--if no encryption protocol is implemented.

TSIG: accomplishes authentication by including cryptographically hashed shared secrets and time stamps in normal DNS packets sent between name servers and clients. Because the time stamps (when the hash operation was performed) are encrypted along with the shared secret, replay attacks are defeated. Some characteristics of TSIG are:

- Utilizes the MD5-HMAC only—simplifies interoperability
- Works well with Dynamic DNS updates
- Authenticates all zone transfer traffic including glue records

TSIG requires manual key distribution: No provision has been made here for distributing the shared secrets; it is expected that a network administrator will statically configure name servers and clients using some out of band mechanism such as sneaker-net until a secure automated mechanism for key distribution is available [RFC 2845].

DNSSec: is based on public key cryptography and digital signatures. Adding data origin authentication and integrity requires no change to the "on-the-wire" DNS protocol beyond the addition of the signature resource type and the key resource type needed for key distribution [RFC 2535].

- Provides for a free-standing PKI that can be used by DNS and other network services
- Don’t have to worry about key distribution
- Additional features come at a cost--added complexity

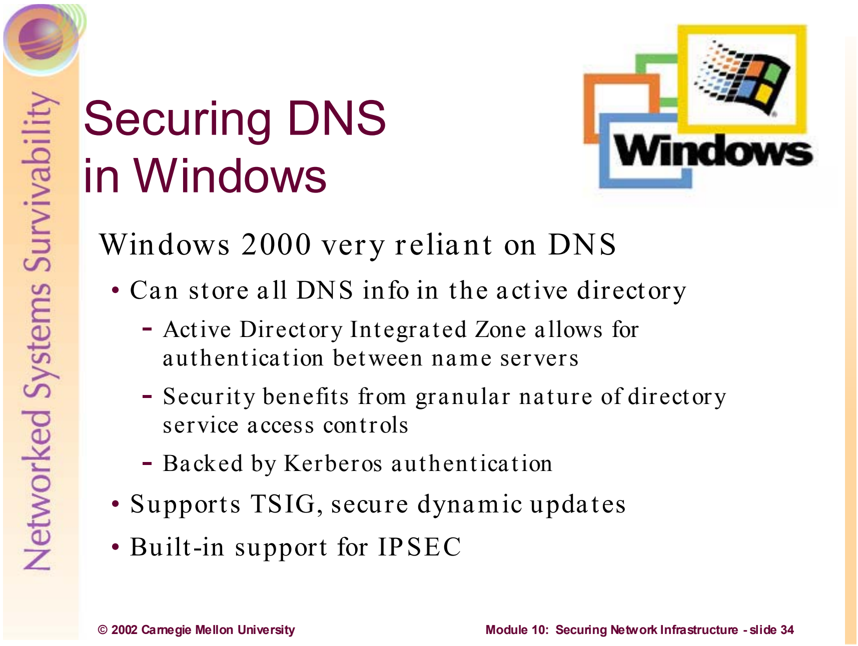
Networked Systems Survivability
Windows

Securing DNS in Windows

Windows 2000 very reliant on DNS

- Can store all DNS info in the active directory
  - Active Directory Integrated Zone allows for authentication between name servers
  - Security benefits from granular nature of directory service access controls
  - Backed by Kerberos authentication
- Supports TSIG, secure dynamic updates
- Built-in support for IPSEC

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 34

Microsoft Windows 2000 networks are designed to be tightly integrated with and dependent upon DNS. For example, when a user logs into a Windows 2000 domain, DNS is required to tell the client what the IP address is of the Windows 2000 domain controller. Windows 2000 networks don't have to use the DNS server that comes bundled with Windows 2000 server, however DNS servers must support Service (SRV) records.

There are advantages to using the Windows 2000 DNS server on Windows 2000 networks:

- Can Integrate the DNS zone database into Active Directory
  - Called an Active Directory Integrated Zone
  - Has the advantages of strong directory service security design which also provides very granular administrative access controls
  - Is replicated to all Windows 2000 Domain Controllers automatically as part of normal directory replication process
  - Has built in authentication (utilizing Kerberos) options when configuring zone transfers. Can select by specific server domain name, or IP address.
- Supports TSIG for native DNS cryptographic authentication
- Supports Secure Dynamic updates
- Can benefit from utilizing IPSec throughout network or just between name servers IPSec is fully integrated into Windows 2000 and it can be implemented easily

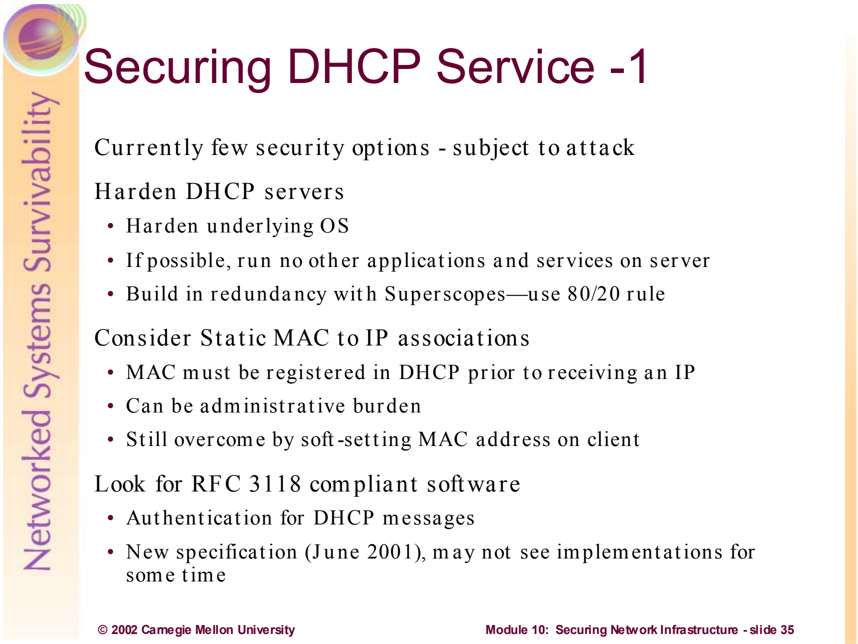

Networked Systems Survivability

Securing DHCP Service -1

Currently few security options - subject to attack

Harden DHCP servers

- Harden underlying OS
- If possible, run no other applications and services on server
- Build in redundancy with Superscopes—use 80/20 rule

Consider Static MAC to IP associations

- MAC must be registered in DHCP prior to receiving an IP
- Can be administrative burden
- Still overcome by soft-setting MAC address on client

Look for RFC 3118 compliant software

- Authentication for DHCP messages
- New specification (June 2001), may not see implementations for some time

© 2002 Carnegie Mellon University   Module 10: Securing Network Infrastructure - slide 35

The Dynamic Host Configuration Protocol (DHCP) supplies IP addresses and other configuration information (like default gateway, DNS servers, etc) to network clients setup to use DHCP. Because DHCP relieves the significant administrative burden of manually configuring network hosts, it is almost universally implemented.

The following is excerpted directly from DHCP's specifying RFC (number 2131):

DHCP is built directly on UDP and IP which are as yet inherently insecure. Furthermore, DHCP is generally intended to make maintenance of remote and/or diskless hosts easier. While perhaps not impossible, configuring such hosts with passwords or keys may be difficult and inconvenient. Therefore, DHCP in its current form is quite insecure.

Unauthorized DHCP servers may be easily set up. Such servers can then send false and potentially disruptive information to clients such as incorrect or duplicate IP addresses, incorrect routing information (including spoof routers, etc.), incorrect domain nameserver addresses (such as spoof nameservers), and so on. Clearly, once this seed information is in place, an attacker can further compromise affected systems. Malicious DHCP clients could masquerade as legitimate clients and retrieve information intended for those legitimate clients. Where dynamic allocation of resources is used, a malicious client could claim all resources for itself, thereby denying resources to legitimate clients.

There are some practices that can help minimize the risk of some of these attacks:

- Secure the operating system that the DHCP application is running on (see harden the name server in previous section on DNS Security)
- Build redundancy by using at least 2 DHCP servers on your network. Set up Superscopes on your servers where 80 percent of a DHCP server's pool of IP addresses are leased locally and the remaining 20 percent are leased from a second DHCP server. This allows DHCP leases to be granted even when one of the servers fails.
- It is possible to associate specific MAC addresses with IP addresses in DHCP. The administrator can manually register the MAC addresses of network computers and only offer IPs and configuration information if the DHCP client's MAC address is previously registered with DHCP. This can be an administrative nightmare, and is certainly far from full proof (due to MAC spoofing and soft-setting) but it does make it a bit harder for intruders to gain access to DHCP services.

- Keep on the lookout for updates to DHCP Software that is RFC 3118 compliant[39]. This is a new specification for authenticating DHCP messages. It should solve many of the current security problems associated with DHCP.
- Don't use DHCP for critical servers on your network—these should use static addressing

[39] http://www.rfc-editor.org/rfc/rfc3118.txt

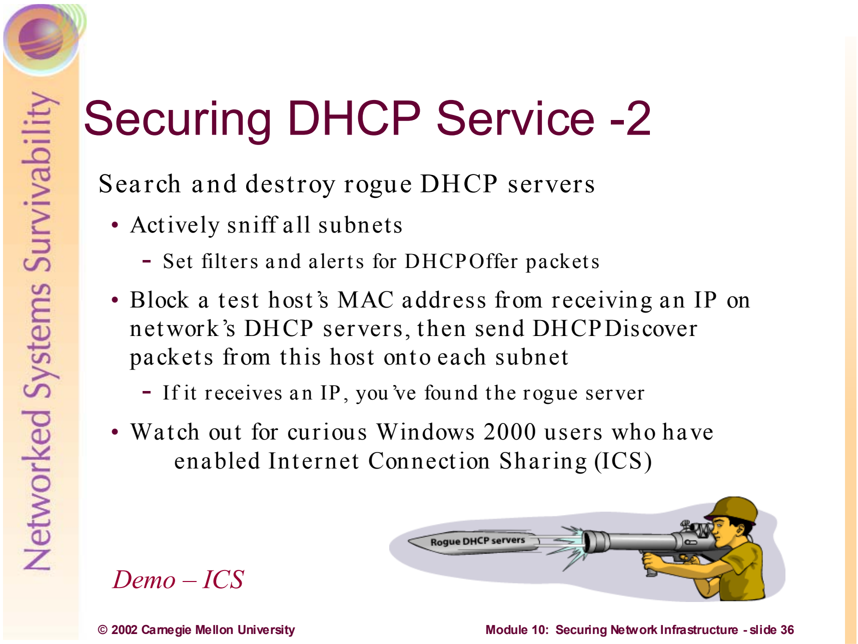

Networked Systems Survivability

Securing DHCP Service -2

Search and destroy rogue DHCP servers

- Actively sniff all subnets
  - Set filters and alerts for DHCPOffer packets
- Block a test host's MAC address from receiving an IP on network's DHCP servers, then send DHCPDiscover packets from this host onto each subnet
  - If it receives an IP, you've found the rogue server
- Watch out for curious Windows 2000 users who have enabled Internet Connection Sharing (ICS)

Demo – ICS

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 36

Rogue (unauthorized) DHCP servers are commonplace in most universities and other network environments. It is very easy for any curious or malicious person to start giving out IP addresses on the local subnet—which is obviously a security problem and can also drive your help desk personnel crazy!

A proactive measure is for administrators to capture packets on all subnets (utilizing software like Sniffer or TCPDump) and to setup filters for DHCP traffic. DHCP runs on UDP ports 67 and 68. Simply set up filters for these ports (UDP port 67 is used by the server, 68 by the client) and specifically set alerts for DHCPOffer packets. DHCPOffer packets are broadcasts from a DHCP server that has just received a DHCPDiscover Packet from a host (who is requesting initial configuration information). A good graphical overview of the four-step DHCP IP address leasing process is available at: http://www.j51.com/~sshay/tcpip/dhcp/dhcp.htm.

Another technique for discovering rogue DHCP servers is to specifically configure your DHCP servers to NOT lease an IP for the MAC address of one (or more) of your administrative hosts. Then configure this host as a DHCP client and wait. If that client receives an IP address, then you know you must have a rogue serve somewhere on that subnet. Tracking down the actual physical location of that system can be somewhat arduous as you'll have to check the switch port to MAC associations on all switches for that subnet (or Vlan).

Windows 2000 Professional comes with a new feature called Internet Connection Sharing (ICS). ICS is designed to share a single connection to an ISP in a small office/home office (SOHO) network. Basically, when an ICS system starts up, Network Address Translation (NAT), DNS proxy, and DHCP services are started automatically. DHCP enabled clients on the local subnet will be leased a private IP address (in the 192.168.0.1/24 range). These clients have their DNS server and default gateway configurations set to use the private IP of the ICS enabled system.

ICS can be enabled with a few mouse clicks (see screen shot below). If a curious network user stumbles across this and decides to activate it, he/she has just enabled a rogue DHCP server on your network that can cause a real mess in a hurry.

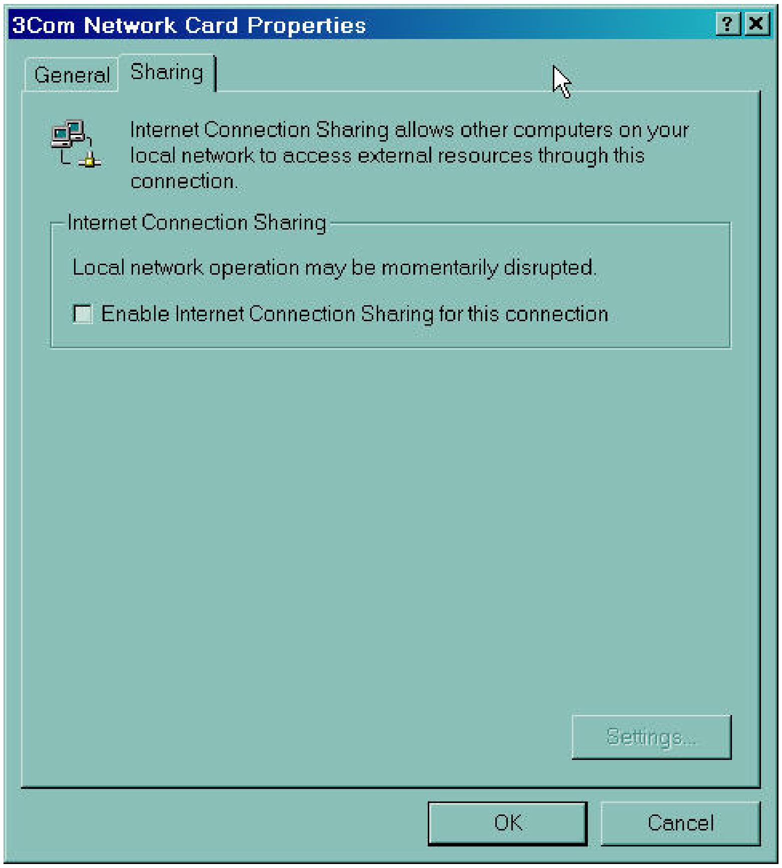

3Com Network Card Properties

General | Sharing

Internet Connection Sharing allows other computers on your local network to access external resources through this connection.

Internet Connection Sharing

Local network operation may be momentarily disrupted.

□ Enable Internet Connection Sharing for this connection

As a best practice, ICS should be disabled via group policy on all Windows 2000 machines in your (non-SOHO) network. Also, use the private address range (set up sniffer filters) as a means of troubleshooting suspected ICS systems.

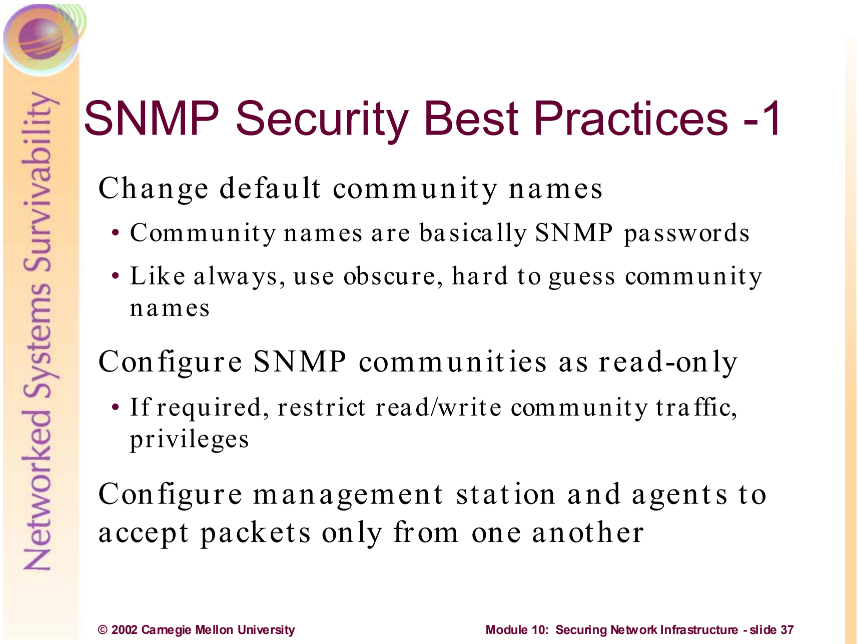

Networked Systems Survivability

SNMP Security Best Practices -1

Change default community names

- Community names are basically SNMP passwords
- Like always, use obscure, hard to guess community names

Configure SNMP communities as read-only

- If required, restrict read/write community traffic, privileges

Configure management station and agents to accept packets only from one another

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 37

Simple Network Management Protocol (SNMP) is widely used by network administrators to monitor and administer all types of computers (e.g., routers, switches, printers). SNMP (version 1)[40] uses an unencrypted "community string" as its only authentication mechanism. Attackers can use this vulnerability in SNMP to possibly gather information from, reconfigure or shut down a computer remotely.

If an attacker can collect SNMP traffic on a network, then he can learn a great deal about the structure of the network as well as the systems and devices attached to it. Disable all SNMP servers on any computer where it is not necessary. However, if SNMP is a requirement, then consider the following. Allow read-only access and not read-write access via SNMP. Do not use standard community strings (e.g., public, private). If possible, only allow a small set of computers access to the SNMP server on the computer [NSA2].

[40] ftp://ftp.isi.edu/in-notes/rfc1157.txt

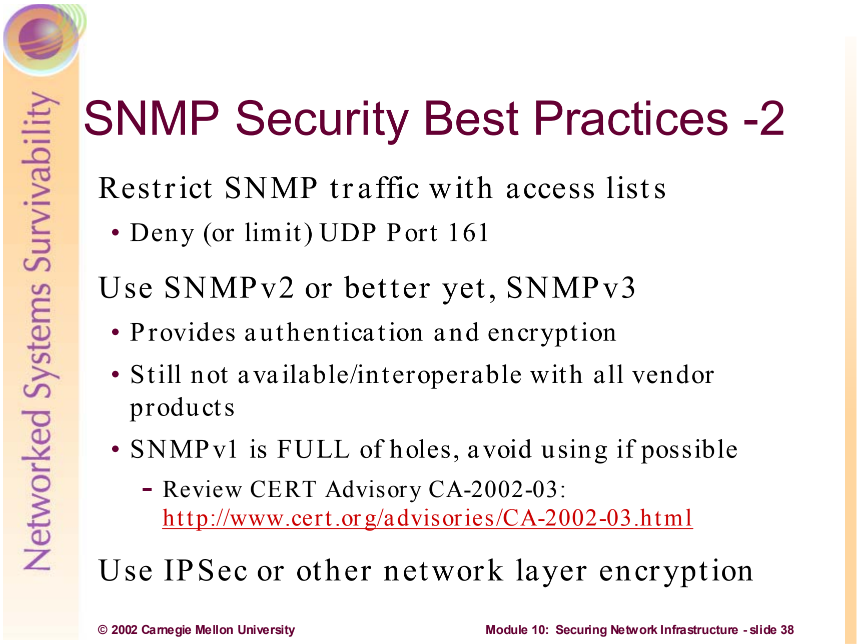
Networked Systems Survivability

SNMP Security Best Practices -2

Restrict SNMP traffic with access lists

- Deny (or limit) UDP Port 161

Use SNMPv2 or better yet, SNMPv3

- Provides authentication and encryption
- Still not available/interoperable with all vendor products
- SNMPv1 is FULL of holes, avoid using if possible
    - Review CERT Advisory CA-2002-03: http://www.cert.org/advisories/CA-2002-03.html

Use IPSec or other network layer encryption

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 38

It is wise to restrict SNMP traffic on your network. Access lists can be configured on routers to do just that. Here is a sample Cisco Access list that blocks SNMP traffic:

```
Router(config)# Access-list 101 deny udp any any eq 161 log
Router(config)# Access-list 101 permit IP any any
```

The first line can be changed to allow traffic between specific hosts only. The above access list (as written) should be applied to interfaces that connect to networks where SNMP is denied.

As mentioned, SNMP community strings are sent on the network in cleartext ASCII. Thus, anyone who has the ability to capture a packet on the network can discover the community string. This may allow unauthorized users to query or modify routers via SNMP. For this reason, using the Cisco IOS command *no snmp-server trap-authentication* may prevent intruders from using trap messages (sent between SNMP managers and agents) to discover community strings.[41]

Because of the inherent security flaws in SNMPv1, it is recommended that either version 2[42] or version 3[43] be implemented. Both of these provide cryptographic authentication and encryption of SNMP traffic, however many network systems still only support SNMPv1. Check the specifications of your equipment carefully and use the most secure version of SNMP where available. CERT advisory 2002-03[44] describes numerous vulnerabilities in SNMPv1 that affects more vendor products than ANY previous advisory. If you must implement SNMPv1, consider implementing IPSec or some other network layer encryption method.

[41] http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs003.htm#xtocid344214

[42] ftp://ftp.isi.edu/in-notes/rfc1446.txt

[43] ftp://ftp.isi.edu/in-notes/rfc2570.txt

[44] http://www.cert.org/advisories/CA-2002-03.html

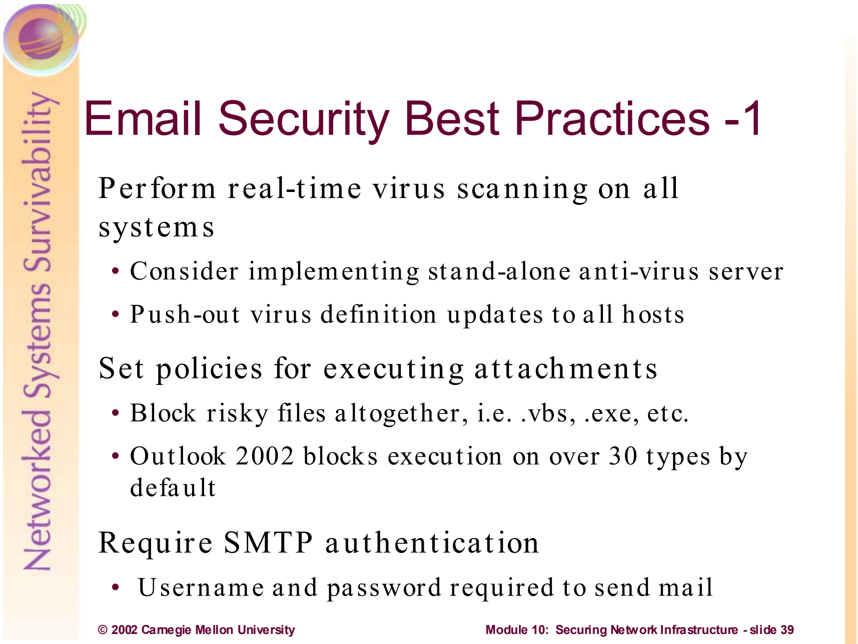

Networked Systems Survivability

Email Security Best Practices -1

Perform real-time virus scanning on all systems

- Consider implementing stand-alone anti-virus server
- Push-out virus definition updates to all hosts

Set policies for executing attachments

- Block risky files altogether, i.e. .vbs, .exe, etc.
- Outlook 2002 blocks execution on over 30 types by default

Require SMTP authentication

- Username and password required to send mail

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 39

Email is the single most utilized network service in the world. As a result, it is critical to almost all organizations. Here are some email security best practices that can enhance survivability:

Implement effective Anti-Virus technology and policy:

Internet viruses and worms have been getting more and more prevalent and destructive. It goes without saying that organizations should implement an aggressive anti-virus strategy. This strategy should be layered, such that virus scanning is done on servers as well as clients.

It is recommended that a separate anti-virus server be implemented. This server can be located just inside the firewall (possibly in an internal DMZ) and should receive inbound email from the external mail forwarder. It should scan all emails for viruses prior to sending them on to the actual mail servers. It should also scan outbound email as well, prior to sending them through the firewall to the mail forwarder. This implementation frees up the mail servers and isolates potentially crippling malicious code.

It is equally important to keep the anti-virus definition files current—on both the servers and the clients. Most enterprise anti-virus software applications have automated update features that can make this process transparent to both the servers and the clients. Implement this--definitely don't rely on your network users to keep their client anti-virus definitions current!

Most virus scanning products function based upon scans for known virus signatures; therefore, they are ineffective against new or uncharacterized attacks. However, they can be effective at preventing reoccurrences of past attacks. Most anti-virus products allow the blocking of mail attachments at the mail server – this may be of value in stopping an outbreak of attachment-based malicious code within an organization in the interim before an update to an antiviral scanning tool's signature file is available. Being able to block attachments would allow basic e-mail connectivity but preclude infection by viruses that use attachments as the transport media such as Nimda and the ILOVEYOU attacks.

There are numerous kinds of executable file attachments that many organizations do not need to routinely distribute via e-mail. If possible, block these at the perimeter as a countermeasure against the malicious code threat. Organizations using Outlook can also block them using Outlook 2002 (blocks over 30 by default) or, for earlier versions of Outlook, by using the appropriate security patches.

The specific file types that should be blocked are:

.bas .hta .msp .url .bat .inf .mst .vb .chm .ins .pif .vbe .cmd .isp .pl .vbs .com .js .reg .ws .cpl .jse .scr .wsc .crt .lnk .sct .wsf .exe .msi .shs .wsh

It may be prudent to add, or delete files from this list depending upon operational realities. For example, it may be practical to block applications within the Microsoft Office family, all of which can contain an executable component. Most notable are Microsoft Access files, which unlike other members of the Office family have no intrinsic protection against malicious macros [NSA2].

Respect the concept of least privilege:

Least privilege is a basic tenet of computer security that basically means "giving a user only those rights that s/he needs to do their job". Executable content runs in the security context on which it was launched – practically speaking, this means in the context of the user launching the code. Good practices include making certain that administrative accounts are kept to a minimum, that administrators use a regular account as much as possible instead of logging in as administrator to do routine things such as reading their mail, and setting resource permissions properly. This will limit the access of any malicious executables that may be inadvertently launched.

Consider implementing SMTP authentication:

This adds one more layer in your overall email security strategy. SMTP authentication simply requires that users input a username and password before any of their mail will be sent. Most client applications have configurable options for this—all of these will store the users' required information and make the process transparent to the user. This doesn't add a great amount of security, it just supports the goal of defense in depth.

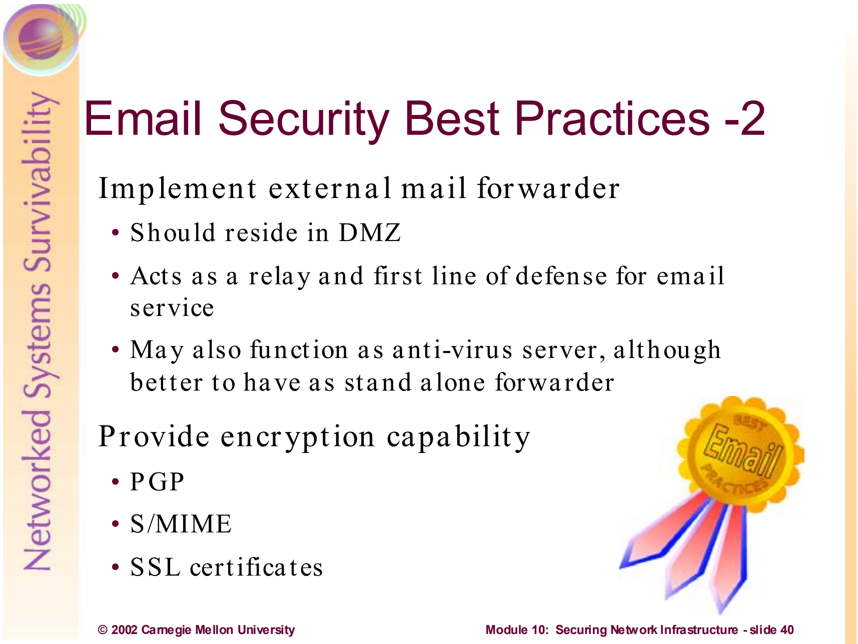

Networked Systems Survivability

Email Security Best Practices -2

Implement external mail forwarder

- Should reside in DMZ
- Acts as a relay and first line of defense for email service
- May also function as anti-virus server, although better to have as stand alone forwarder

Provide encryption capability

- PGP
- S/MIME
- SSL certificates

BEST
Email
PRACTICES

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 40

As mentioned previously, implementing a mail forwarder will enhance email security.

There are several risks associated with the receipt of e-mail from potentially untrusted entities outside the site. Chief among these concerns are attacks against the recipient e-mail server itself. Examples of this include attempts to exploit buffer overflows and content driven attacks in the form of malicious code. The mail forwarder is simply a mail server that forwards e-mail messages intended for internal users to the internal mail server (or anti-virus server) and accepts mail destined for the external network for delivery. As the mail forwarder is the only mail server that is exposed to the external network, the risks associated with e-mail are reduced by precluding direct access to the internal mail server. This mail forwarder, as should be the case for all servers in the DMZ, must not be a member of any internal domain. This will limit the damage that could result from its compromise. Content checking, initial virus scanning, and filtering can optionally be performed here to further guard against malicious code.

The Internal Mail Server is utilized by users within the site to both send and receive mail. It should be configured to send all outgoing mail to the mail forwarder (or anti-virus server). The risks associated with the internal mail server are similar to that described for mail forwarder. Again, content checking and filtering capabilities are critical countermeasures. Additionally, a large concern is preserving data confidentially by utilizing user authentication and access control mechanisms to limit users to content for which they are authorized access (e.g., their own mailbox). Data encryption can also be utilized for sensitive data with the common options including the S/MIME standard for reader-to-writer data protection or SSL for protecting data in transmission between the client and server. PGP is also a widely used method for encrypting email. Many popular Email clients (i.e., Outlook, Mulberry, etc) have PGP plug-ins, that make it simple to send/receive encrypted messages [NSA3].

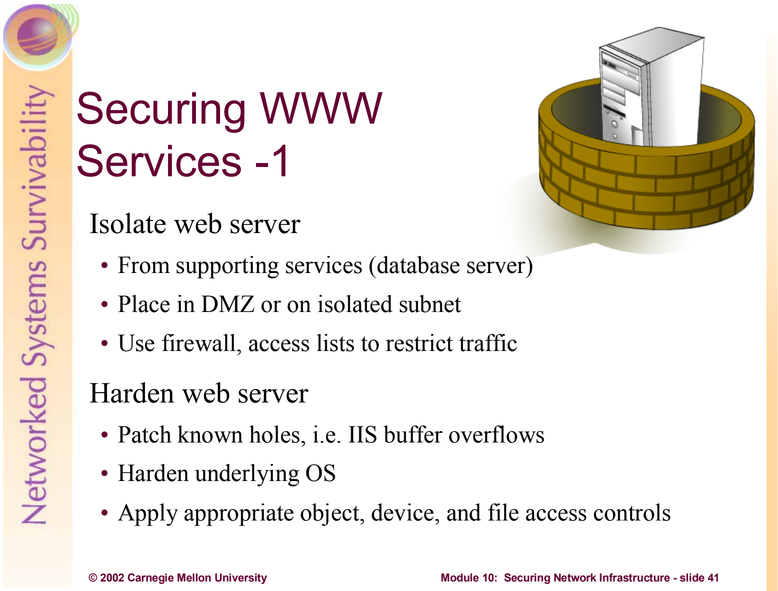

Networked Systems Survivability

Securing WWW Services -1

Isolate web server

- From supporting services (database server)
- Place in DMZ or on isolated subnet
- Use firewall, access lists to restrict traffic

Harden web server

- Patch known holes, i.e. IIS buffer overflows
- Harden underlying OS
- Apply appropriate object, device, and file access controls

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 41

A public Web server host is a computer intended for public access. This means that there will be many people who will access the host (and its stored information) from locations all over the world. Regardless of how well the host computer and its application software are configured, there is always the chance that someone will discover a new vulnerability, exploit it, and gain unauthorized access to the Web server host (e.g., via a user account or a privileged account on a host with a multiuser operating system). If that occurs, you need to prevent these subsequent events, if possible:

- The intruder is able to observe or capture network traffic that is flowing between internal hosts. Such traffic might include authentication information, proprietary business information, personnel data, and many other kinds of sensitive data.
- The intruder is able to get to internal hosts, or to obtain detailed information about them.

To guard against these two threats, the public Web server host must be isolated from your internal network and its traffic, as well as from any public network.

Place the Web server on a subnet isolated from public and internal networks (DMZ). By doing so, network traffic destined for the Web server subnet can be better monitored and controlled. This aids in configuring any firewall or router used to protect access to the subnet as well as detecting attacks and attempted intrusions. It also precludes the capture of internal traffic (accessible to all connected computers when using broadcast media such as Ethernet) by any intruder who gains access to your Web server.

Use firewall technology to restrict traffic between a public network and the Web server, and between the Web server and internal networks. The use of firewall technology[1] (including packet filtering implemented by a router) effectively restricts the traffic between all computers connected to the firewall in accordance with your security policy. Setting up firewall technology precludes many possible attacks but still allows anyone to access your public Web server content. A Web server typically accepts TCP connections on port 80/tcp (http), the standard port. In general, no other connections from the public network to the Web server should be permitted. However, if the Web server supports SSL - protected connections, port 443/tcp (https) should be permitted.

You need to establish filtering rules that block TCP connections originating from the Web server, as a Web server typically does not depend on other services on the public network. In general, all UDP and ICMP traffic should be blocked. However, depending on the Web server configuration, you may need to permit limited connections to UDP-based services such as DNS for host name lookup (permit port 53/udp). One recommended configuration is shown below [45]

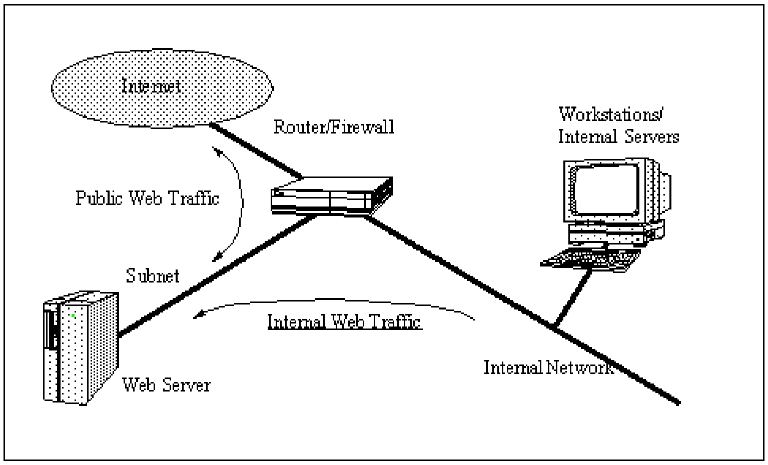

Most Web server host operating systems provide the capability to specify access privileges individually for files, devices, and other data or code objects stored on that host. Any information that your Web server can access using these controls can potentially be distributed to all users accessing your public Web site. Your Web server software is likely to provide additional object, device, and file access controls specific to its operation. You need to consider how best to configure these access controls to protect information stored on your public Web server from two perspectives:

- to limit the access of your Web server software
- to apply access controls specific to the Web server where more detailed levels of access control are required

The proper setting of access controls can help prevent the disclosure of sensitive or restricted information that is not intended for public dissemination. In addition, access controls can be used to limit resource use in the event of a denial-of-service (DoS) attack against your public Web site.[46]

Harden OS and patch vulnerabilities and holes as mentioned earlier in this module.

[45] http://www.cert.org/security-improvement/practices/p075.html
[46] http://www.cert.org/security-improvement/practices/p076.html

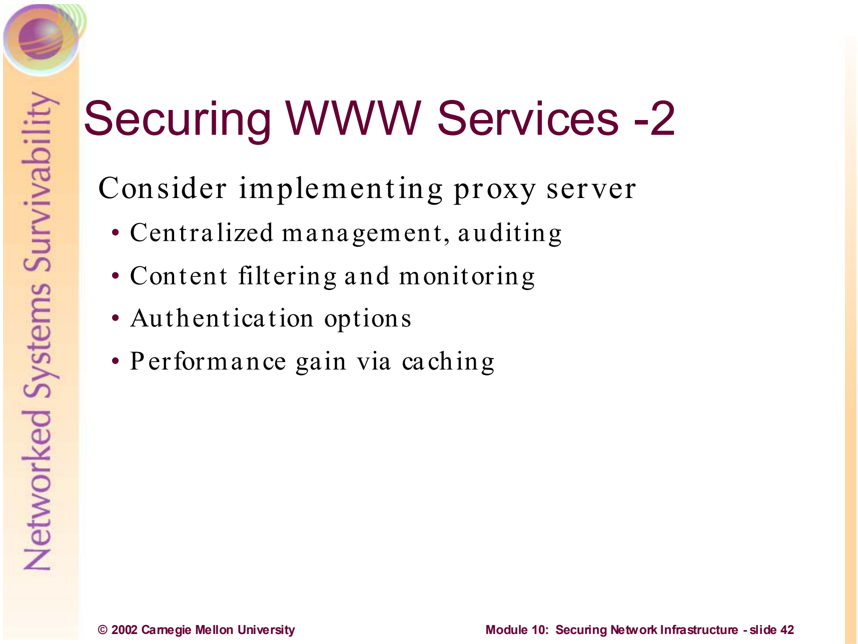

Networked Systems Survivability

Securing WWW Services -2

Consider implementing proxy server

- Centralized management, auditing
- Content filtering and monitoring
- Authentication options
- Performance gain via caching

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 42

HTTP proxy servers can provide a great deal of flexibility when it comes to securing the client side of WWW services. Because of this, it is recommended that they be instituted in larger networks.

HTTP proxy servers provide a centralized management capability for WWW traffic. It is analogous to a valve on a water line. The valve allows you to control how much water gets through, and with some creating plumbing, allows you to select where the water is from and where it is going. They also provide a central location for instituting content filtering and monitoring.

Some proxy servers allow you to authenticate users prior to gaining access to the WWW service. Thereby different groups could have different monitoring and filtering applied.

Caching technology is becoming very mature, and it can significantly improve the efficiency of bandwidth utilization. The “surfing experience” will seem faster for your users, too!

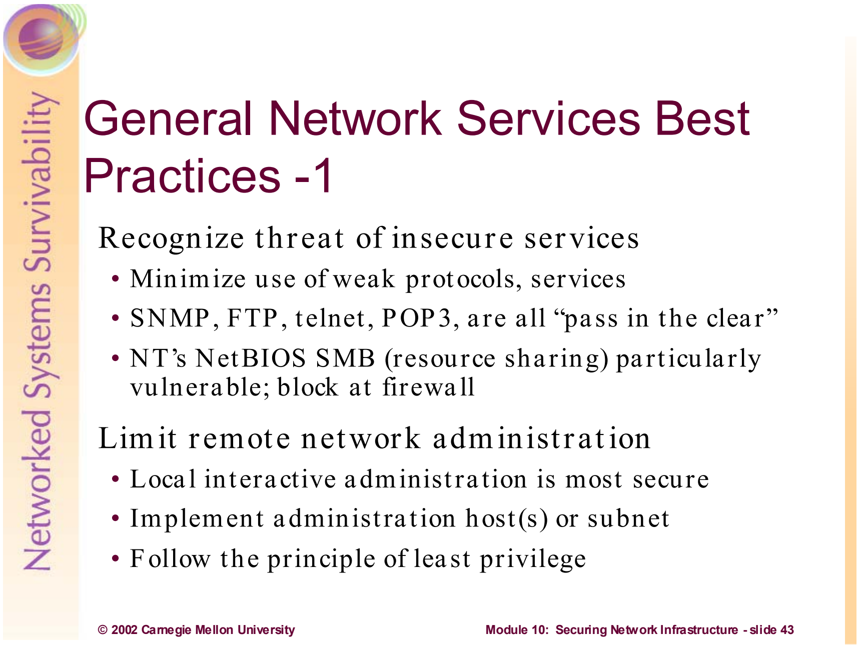

Networked Systems Survivability

General Network Services Best Practices -1

Recognize threat of insecure services

- Minimize use of weak protocols, services
- SNMP, FTP, telnet, POP3, are all “pass in the clear”
- NT’s NetBIOS SMB (resource sharing) particularly vulnerable; block at firewall

Limit remote network administration

- Local interactive administration is most secure
- Implement administration host(s) or subnet
- Follow the principle of least privilege

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 43

This is a summary slide for the content covered previously.

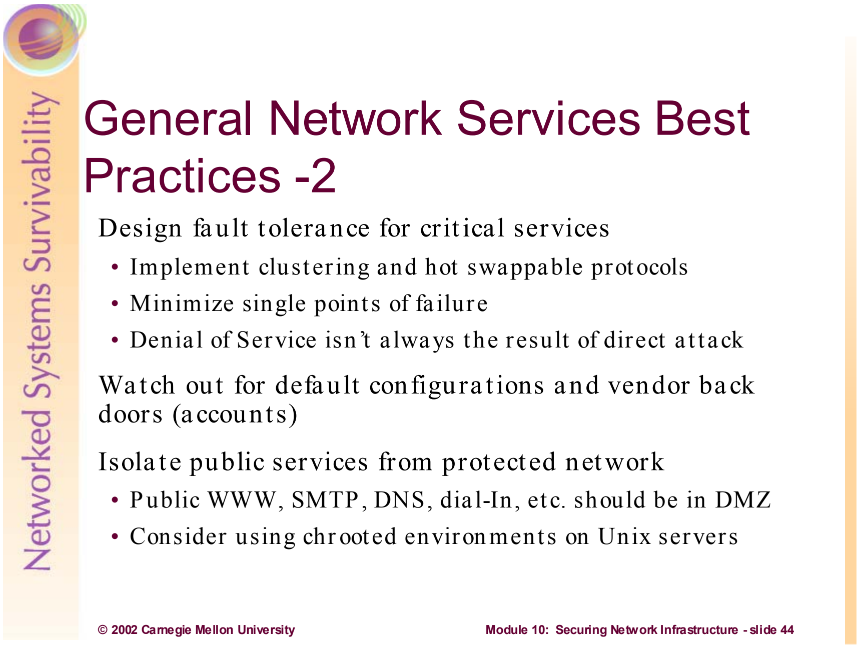

Networked Systems Survivability

General Network Services Best Practices -2

Design fault tolerance for critical services

- Implement clustering and hot swappable protocols
- Minimize single points of failure
- Denial of Service isn't always the result of direct attack

Watch out for default configurations and vendor back doors (accounts)

Isolate public services from protected network

- Public WWW, SMTP, DNS, dial-In, etc. should be in DMZ
- Consider using chrooted environments on Unix servers

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 44

Again this is another summary slide.

See a comprehensive list of vendor back doors and default passwords at: http://www.phenoelit.de/dpl/dpl.html

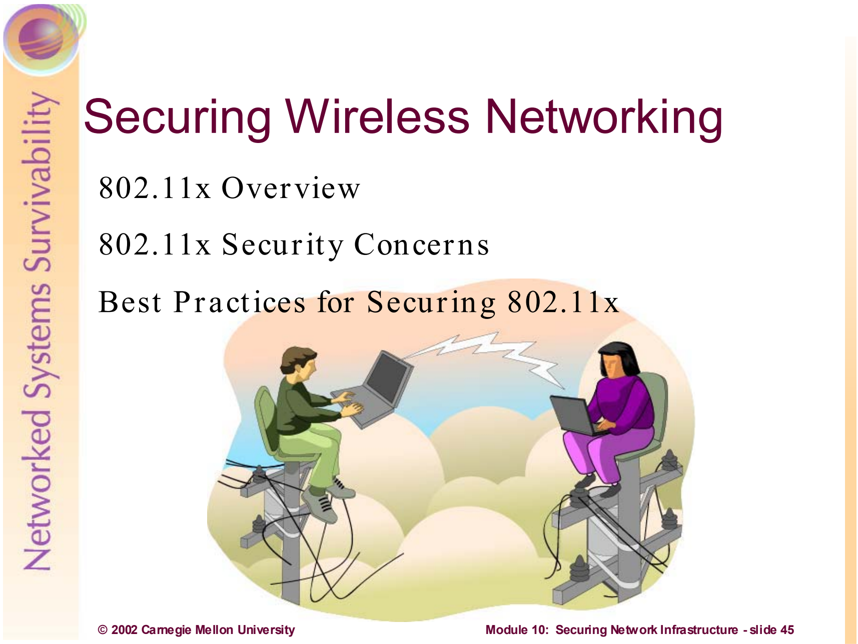

Networked Systems Survivability

Securing Wireless Networking

802.11x Overview

802.11x Security Concerns

Best Practices for Securing 802.11x

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 45

Wireless networks are booming. It seems that nearly all universities, businesses, commercial airports, and all levels of government are going wireless. Walt Disney World has a vast wireless network deployed utilizing the IEEE's 802.11b standard.[47] Carnegie Mellon University has the most extensive wireless network of any university in the country.[48] I and almost all of my co-workers have wireless networks in our homes. I've written this module (on my laptop) in my basement, living room, and bedroom and have had access to my broadband Internet connection all along. Why is this proliferation happening? Because wireless networking is really cool—we'll get more technical soon enough! It's affordable, easy to deploy, and very convenient.

Those of us who've been around computer networking for a while have seen this kind of booming technology before. It has become a predictable and recurring scenario when it comes to emerging technologies--security is almost never considered and is usually immature at best. Wireless networking is no different. The 802.11x standard is insecure, however there are some things you can do to lessen the risk.

[47] http://www.computerworld.com/storyba/0,4125,NAV47_STO65816,00.html
[48] http://www.yil.com/wiredcolleges/top100chart.pdf

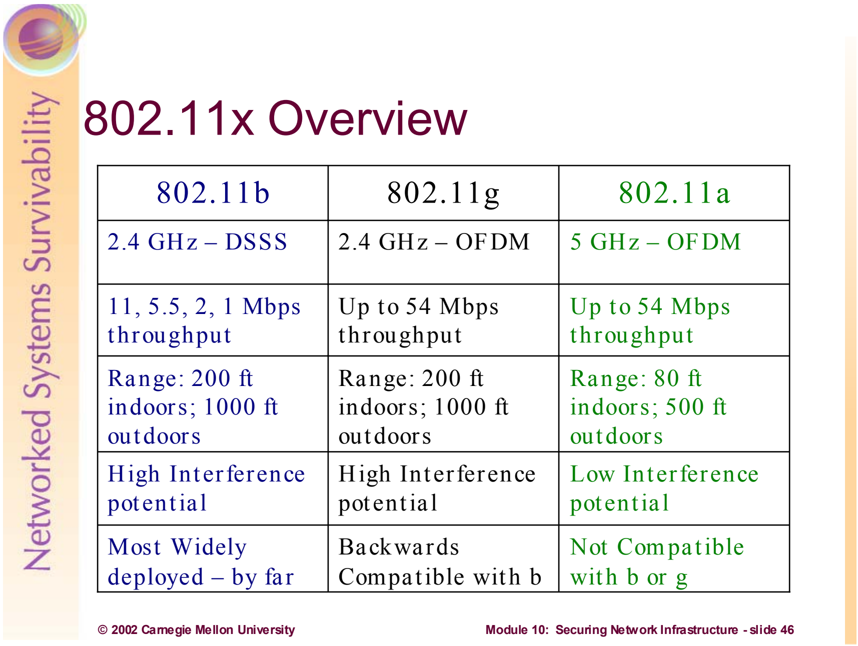
Networked Systems Survivability

802.11x Overview

802.11b	802.11g	802.11a
2.4 GHz – DSSS	2.4 GHz – OFDM	5 GHz – OFDM
11, 5.5, 2, 1 Mbps throughput	Up to 54 Mbps throughput	Up to 54 Mbps throughput
Range: 200 ft indoors; 1000 ft outdoors	Range: 200 ft indoors; 1000 ft outdoors	Range: 80 ft indoors; 500 ft outdoors
High Interference potential	High Interference potential	Low Interference potential
Most Widely deployed – by far	Backwards Compatible with b	Not Compatible with b or g

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 46

As of the time of this writing, the IEEE 802.11b standard is still the predominately deployed standard for wireless networking. Its specifications are constrained to the physical and data link layers of the OSI reference Model[49]. The "b" in 802.11b merely specifies the high-data rate implementation of the original 802.11[50] standard, which topped out at 2 Mbps. 802.11b[51] utilizes a 2.4 GHz radio frequency (RF) carrier and some refinements in media access control to extend the bandwidth to a maximum of 11Mbps over Ethernet. To help attain this gain in bandwidth, it specifies the use of Direct Sequence Spread Spectrum (DSSS)[52] which is a means of multiplexing many signals onto one RF carrier.

Under 802.11b, devices communicate at a speed of 11 Mbps whenever possible. If signal strength or interference is disrupting data, the devices will drop back to 5.5 Mbps, then 2 Mbps and finally down to 1 Mbps. Though it may occasionally slow down, this keeps the network stable and very reliable.

The span of coverage utilizing 1 access point (no-roaming) varies widely. The diameter of coverage when you're outdoors (with no obstructions between your wireless card and the AP) can extend to more than 1000 feet. Indoors coverage is much less—usually less than 100 feet in most buildings. Antennas can be used to boost the signal significantly--one very effective 802.11b antenna can be made for less than $10 utilizing a Pringles potato chip can![53] The signal is subject to interference from other RF signals, physical obstructions like buildings, and weather conditions. Interestingly, this adds a mild layer of security because it limits the range of the data signal. Rain interferes a great deal with the 802.11b signal—so you are probably more secure if your wireless network is located in Seattle!

Newer standards have been researched and approved by the IEEE in the recent past for wireless networking. These include 802.11a and 802.11g, both of which are advertised to transmit data at 54 Mbps. They differ from the 802.11b standard only in the media (radio frequencies used and multiplexing scheme). The effective data rate of both of these standards approaches 25 Mbps, which is about 6-8 times that of that of 802.11b.

[49] http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/introint.htm#xtocid5

[50] http://standards.ieee.org/reading/ieee/std/lanman/802.11-1999.pdf

[51] http://standards.ieee.org/reading/ieee/std/lanman/802.11b-1999.pdf

[52] http://www.webopedia.com/TERM/D/DSSS.html

[53] http://www.oreillynet.com/cs/weblog/view/wlg/448

The 802.11a standard uses a completely separate portion of the radio spectrum (5 GHz) than do 802.11b and 802.11g (both at 2.4 GHz, part of the Unlicensed National Information Infrastructure (UNII) band). Because of this change in radio frequency, 802.11a systems are incompatible with the other standards. Also, because of the change in frequency range, 802.11a is more susceptible to attenuation due to weather or other environmental factors, as well as attenuation due to walls or other obstacles when operated inside a building. However, operating in a non-UNII band reduces the potential for interference for the 802.11a standard equipment.

The 802.11g standard, as well as 802.11a, changes from the 802.11b standard in its multiplexing scheme, moving from Direct Sequence Spread Spectrum (DSSS) to Orthogonal Frequency Division Multiplexing (OFDM) to allow for higher data rates. OFDM works by splitting the radio signal into multiple smaller sub-signals that are then transmitted simultaneously at different frequencies to the receiver. OFDM reduces the amount of crosstalk in signal transmissions.

Despite the advances in throughput, no changes made to the new standards address any of the security issues which plagued 802.11b. In fact, the underlying protocols are exactly the same as those used in 802.11b.

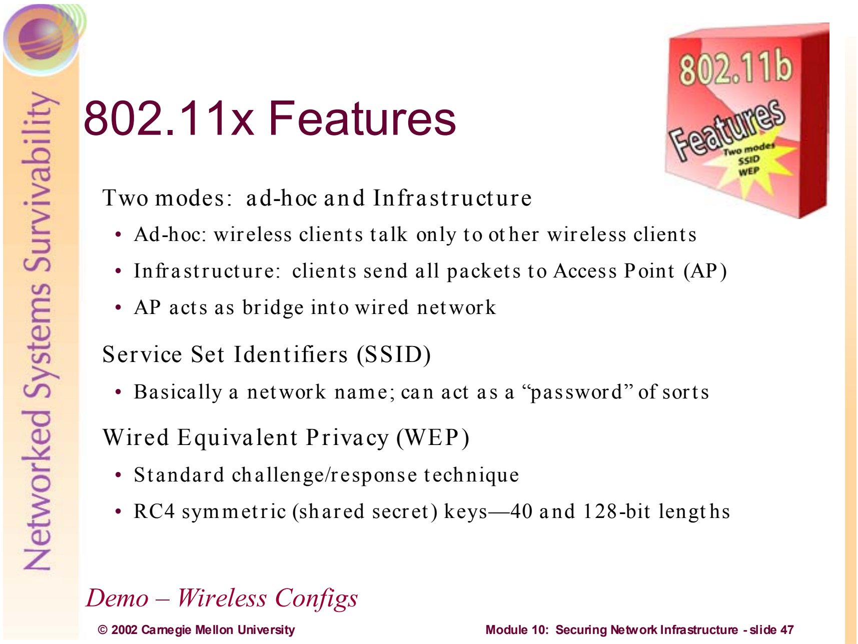
Networked Systems Survivability
802.11b
Features
Two modes
SSID
WEP

802.11x Features

Two modes: ad-hoc and Infrastructure

- Ad-hoc: wireless clients talk only to other wireless clients
- Infrastructure: clients send all packets to Access Point (AP)
- AP acts as bridge into wired network

Service Set Identifiers (SSID)

- Basically a network name; can act as a "password" of sorts

Wired Equivalent Privacy (WEP)

- Standard challenge/response technique
- RC4 symmetric (shared secret) keys—40 and 128-bit lengths

Demo – Wireless Configs

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 47

802.11x has two different operating modes: ad-hoc and infrastructure.

In ad hoc mode: A network is composed solely of stations within mutual communication range of each other via the wireless medium. Basically, users configure their wireless cards to talk to eachother. This is done by implementing the same SSID and normal network protocol (TCP/IP) settings. If any external communications is required, one of the ad hoc systems will have to be configured as a gateway and then route packets to another network.

In infrastructure mode, each client sends all of its communications to a central station, or access point (AP). The access point acts as an Ethernet bridge and forwards the communications onto the appropriate network–either the wired network, or the wireless network.

The specification utilizes Service Set Identifiers (SSID) to identify individual wireless networks. An SSID is just a string of characters (i.e. Cisco Aironet products use *tsunami* as the default SSID) that are chosen and input by an administrator. They function like SNMP community names in a sense—each participating station must be configured to use the same SSID in order to communicate. Most wireless APs available now have the capability for only allowing access if clients are configured with the correct WEP—acts as a shared secret in a sense. Although generally not a well kept secret.

Wired Equivalent Privacy (WEP) is defined as protecting authorized users of a wireless LAN from casual eavesdropping. This service is intended to provide functionality for the wireless LAN equivalent to that provided by the physical security attributes inherent to a wired medium. Wireless transmissions are easier to intercept than transmissions over wired networks. The 802.11 standard currently specifies the WEP security protocol to provide encrypted communication between the client and an AP. WEP employs the symmetric key encryption algorithm, Ron's Code 4 Pseudo Random Number Generator (RC4 PRNG).

Under WEP, all clients and APs on a wireless network use the same key (shared secret) to encrypt and decrypt data. The key resides in the client computer and in each AP on the network. The 802.11 standard does not specify a key management protocol, so all WEP keys on a network must be managed manually. Support for WEP is standard on most current 802.11 cards and APs. WEP security is not available in ad hoc (or peer-to-peer) 802.11 networks that do not use APs.

WEP specifies the use of a 40-bit encryption key and there are also implementations of 104-bit keys. The encryption key is concatenated with a 24-bit "initialization vector," resulting in a 64- or 128-bit key. This key is input into a pseudorandom number generator. The resulting sequence is used to encrypt the data to be transmitted.

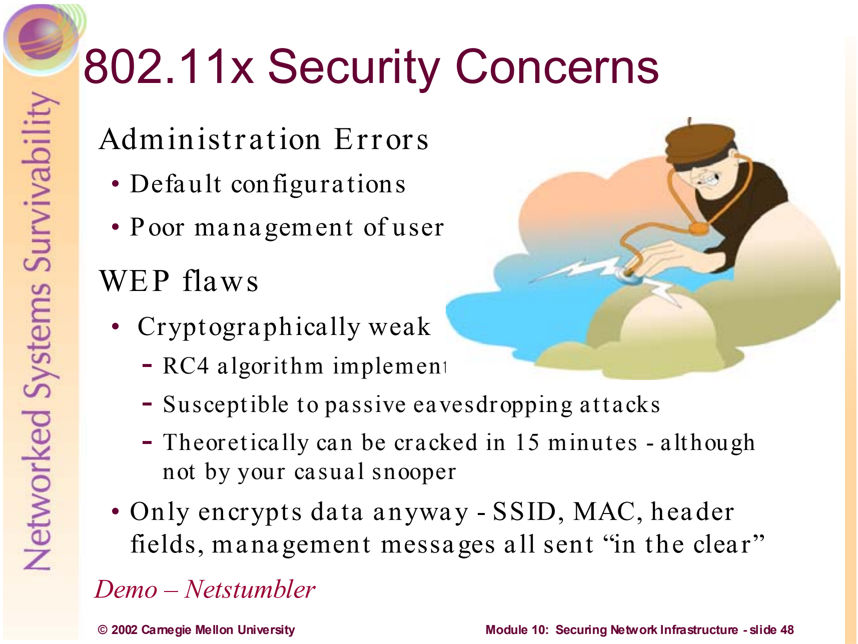

Networked Systems Survivability

802.11x Security Concerns

Administration Errors

- Default configurations
- Poor management of user

WEP flaws

- Cryptographically weak
  - RC4 algorithm implemen
  - Susceptible to passive eavesdropping attacks
  - Theoretically can be cracked in 15 minutes - although not by your casual snooper
- Only encrypts data anyway - SSID, MAC, header fields, management messages all sent "in the clear"

Demo – Netstumbler

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 48

It is imperative that administrators change all of the default configurations that are set by nearly all wireless equipment manufacturers. These defaults include SSIDs, WEP keys, access point management passwords and IP addresses, etc. Anyone can easily discover all of these default settings by surfing to the vendor's web site and looking at the equipment documentation.

Additionally, administrators should require some form of user registration prior to gaining access to the wireless network. This is provides for accountability. Wired network authentication should be used for this (i.e. Kerberos or NTLM)—possibly registering user's wireless MAC.

There have been many widely publicized weakness uncovered with regards to WEP. WEP suffers from two critical flaws: vulnerable encryption and a lack of key management. The most devastating attack was theorized by Scott Fluhrer from Cisco Systems, and Itsik Mantin and Adi Shamir from The Weizmann Institute of Science in Israel—three renowned cryptographers. Researchers from AT&T and Rice University successfully implemented this attack and were able to crack standard WEP in about 15 minutes.[54]

What it boils down to is the initialization vector that goes on the front of the secret key used to generate the ciphertext by RC4. By having a known bit pattern that is prepended on the key, it leads to weak keys that will generate known ciphertext output from the RC4 engine. That allows the attacker to go back and decipher what the secret key is that he doesn't know. The initialization vector is simply the first 24 bits of the shared secret. Also, WEP does not encrypt management frames sent between the AP and wireless clients. As a result, SSIDs, MAC addresses, and other header fields are sent in plaintext. This can easily be verified by downloading a free wireless sniffer (like netstumbler[55]) and taking a war driving expedition through your favorite industrial park.[56]

[54] http://www.cs.rice.edu/~astubble/wep/wep_attack.pdf

[55] http://www.netstumbler.com/

[56] http://www.infosecuritymag.com/2002/jan/columns_note.shtml

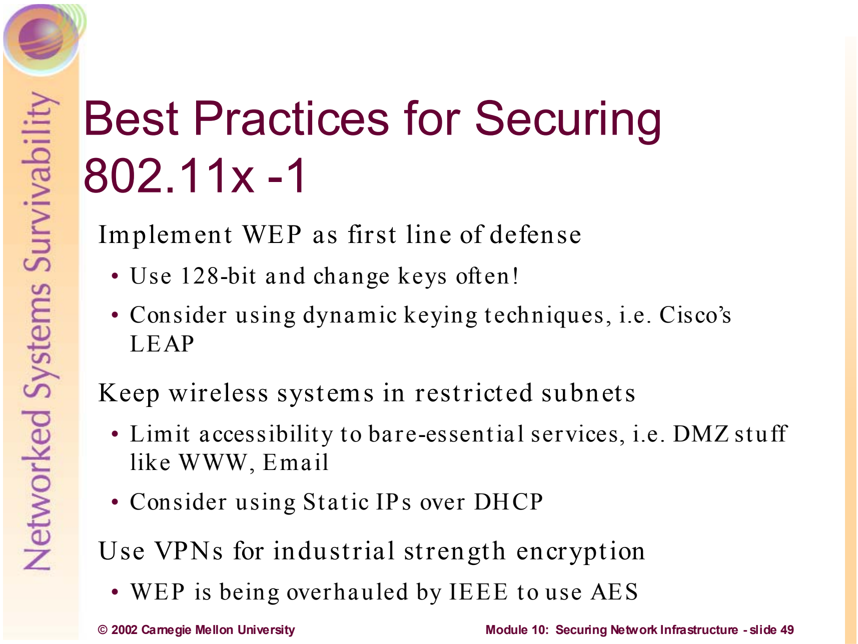
Networked Systems Survivability

Best Practices for Securing 802.11x -1

Implement WEP as first line of defense

- Use 128-bit and change keys often!
- Consider using dynamic keying techniques, i.e. Cisco's LEAP

Keep wireless systems in restricted subnets

- Limit accessibility to bare-essential services, i.e. DMZ stuff like WWW, Email
- Consider using Static IPs over DHCP

Use VPNs for industrial strength encryption

- WEP is being overhauled by IEEE to use AES

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 49

Despite the flaws in WEP, it is certainly better than having no encryption at all. Therefore, it is highly recommend that it be used. Use the strongest WEP keys available and change your keys regularly—every 30 days in a good benchmark. Many vendors have implemented dynamic key management systems (Cisco offers one called LEAP that utilizes the Extensible Authentication Protocol) that automatically change the shared keys at short intervals. This may be a good approach—especially if your wireless equipment vendor is supporting free implementation

Another mitigation technique is to limit the privileges of wireless users. Many users just need access to email and the Internet. This can be done through access controls and firewalls. Segregate the wireless part of your IP network into separate subnets. This makes applying access controls and monitoring easier.

Utilizing statically assigned IPs over DHCP is another means of limiting access to the wireless network. Attempt to obscure the range of valid host IP addresses on your wireless subnets—to make it harder for an intruder to gain access.

A more comprehensive solution to most of the wireless security concerns is to implement virtual private networking (VPN) for ALL of you wireless components. An excellent article from Dell describes how VPNs can help solve your wireless security concerns.[57] See module 12 "Securing Remote Access" for further information on VPNs.

The IEEE 802.11 working group i has been working on a replacement for WEP. The Advanced Encryption Standard has been selected as the basis for this replacement dubbed the Temporal Key Integrity Protocol (TKIP) alongside 802.1x authentication.[58]

[57] http://www.dell.com/us/en/gen/topics/vectors_2001-wireless_security.htm
[58] http://grouper.ieee.org/groups/802/11/

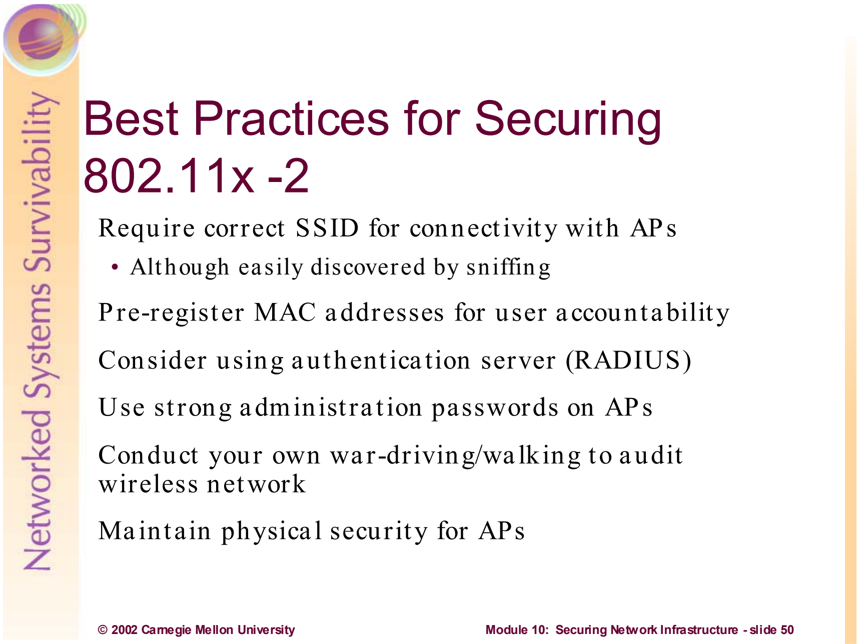
Networked Systems Survivability

Best Practices for Securing 802.11x -2

Require correct SSID for connectivity with APs

- Although easily discovered by sniffing

Pre-register MAC addresses for user accountability

Consider using authentication server (RADIUS)

Use strong administration passwords on APs

Conduct your own war-driving/walking to audit wireless network

Maintain physical security for APs

© 2002 Carnegie Mellon University Module 10: Securing Network Infrastructure - slide 50

Because a client computer must present the correct SSID to access the AP, the SSID acts as a simple password and, thus, provides a measure of security. However, this minimal security is compromised if the AP is configured to "broadcast" its SSID. When this broadcast feature is enabled, any client computer that is not configured with a specific SSID is allowed to receive the SSID and access the AP. In addition, because users typically configure their own client systems with the appropriate SSIDs, they are widely known and easily shared. (It is strongly recommended that APs be configured with broadcast mode disabled.)

While an AP or group of APs can be identified by an SSID, a client computer can be identified by the unique MAC address of its 802.11 network card. To increase the security of an 802.11 network, each AP can be programmed with a list of MAC addresses associated with the client computers allowed to access the AP. If a client's MAC address is not included in this list, the client is not allowed to associate with the AP. MAC address filtering (along with SSIDs) provides improved security, but is best suited to small networks where the MAC address list can be efficiently managed. Each AP must be manually programmed with a list of MAC addresses, and the list must be kept up-to-date. This administrative overhead limits the scalability of this approach.

Utilizing some form of remote user authentication for wireless clients is an excellent means of enhancing security. The AP would pass a user’s credentials on to an authentication server (RADIUS servers are widely implemented) and would only permit connectivity if the user is authenticated.

Additionally, management passwords should follow the same administrative policies as other networking devices (see earlier discussion on passwords under Harden the Router).

It’s a good idea to proactively monitor your own wireless network—go for a war-drive or walk yourself—and while your at it, make sure all of your APs are physically secured.

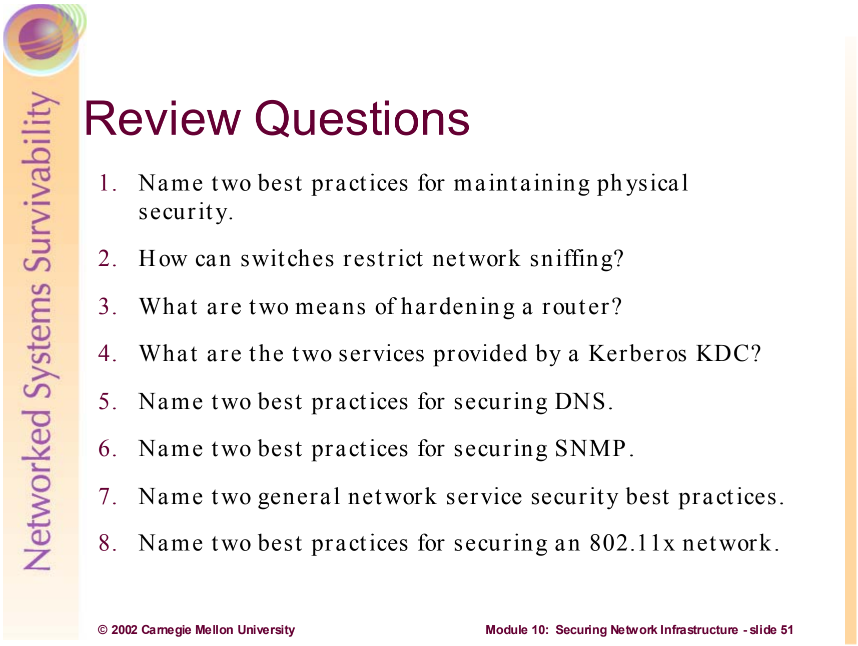

Networked Systems Survivability

Review Questions

1. Name two best practices for maintaining physical security.
2. How can switches restrict network sniffing?
3. What are two means of hardening a router?
4. What are the two services provided by a Kerberos KDC?
5. Name two best practices for securing DNS.
6. Name two best practices for securing SNMP.
7. Name two general network service security best practices.
8. Name two best practices for securing an 802.11x network.

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 51

1. Name two best practices for maintaining physical security.

   Answer: Secure the media and connectivity access points (wall jacks)

2. How can switches restrict network sniffing?

   Answer: Isolate collision domains and broadcast domains with VLans

3. What are two means of Hardening a Router?

   Answer: Minimize use of risky services and features; secure management access

4. What are the two services provided by a Kerberos KDC?

   Answer: Authentication Service and Ticket Granting Service

5. Name two best practices for securing DNS.

   Answer: Implement Split Architecture and implement TSIG or DNSSec

6. Name two best practices for securing SNMP.

   Answer: Change default community strings and restrict read-write access

7. Name two Email Security Best Practices.

   Answer: Actively scan for viruses and implement a mail forwarder

8. Name two best practices for securing an 802.11x network.

   Answer: Implement WEP and change keys often; use VPN technology for all wireless access

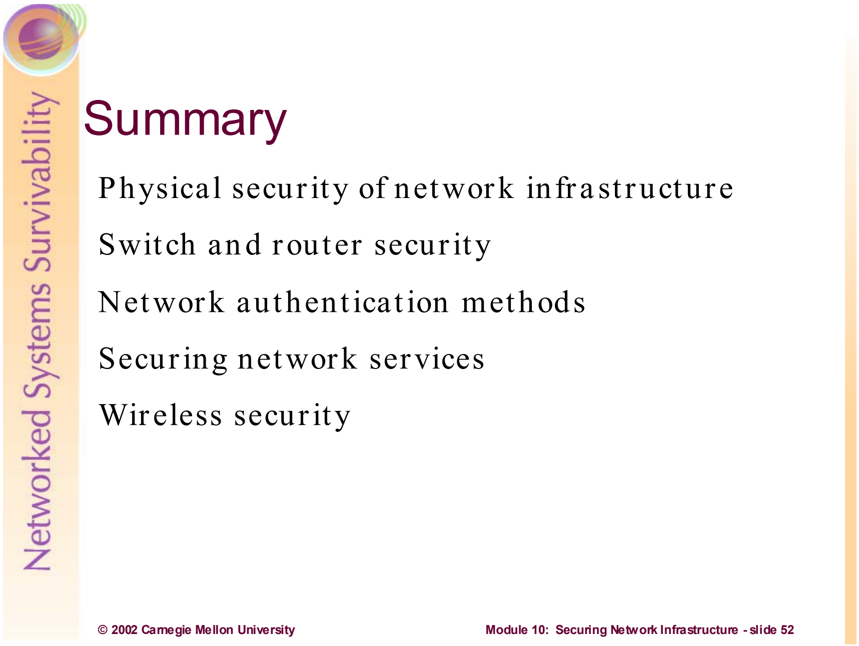

Networked Systems Survivability

Summary

Physical security of network infrastructure

Switch and router security

Network authentication methods

Securing network services

Wireless security

© 2002 Carnegie Mellon University

Module 10: Securing Network Infrastructure - slide 52

References:

[NSA1] National Security Agency; *Router Security Configuration Guide,* Version: 1.0j November 21, 2001. Available at: http://nsa1.www.conxion.com/

[NSA2] National Security Agency; *60 Minute Network Security Guide,* Version: 1.0 October 16, 2001. Available at: http://nsa1.www.conxion.com/support/guides/sd-7.pdf

[NSA3] National Security Agency; *Microsoft Windows 2000 Network Architecture Guide,* Version: 1.0 October 20, 2000. Available at: http://nsa1.www.conxion.com/win2k/guides/w2k-1.pdf

[NCES] National Center for Education Statistics, U.S. Department of Education; *Safeguarding Your Technology*. Available at: http://nces.ed.gov/pubs98/safetech/index.html

[Tyson] Jeff Tyson, HowStuffWorks.Com; *How LAN Switches Work*. Available at: http://www.howstuffworks.com/lan-switch.htm/printable